

Run JAVA Code in C/C++!

C/C++ Users Journal

C/C++ Users Journaltm

Advanced Solutions for C/C++ Programmers

STL

- Cleaner Code with STL's *map*
- Two Important STL Books Compared
- STL Help on the Web

Simple-as-a-Stopwatch Timing Class

Finding the Right Screen Library

Keep Floating-Point Sums Precise

P. J. Plauger

STL's Treasure Trove of Algorithms

Dan Saks

Making C++ Grammar More Intuitive

Pete Becker

When to Break the Rules of OOP

Bobby Schmidt

The Winding Road from C to C++

\$4.95 U.S. Canada \$5.95



September

1996

Vol. 14, No. 9



Communications?

Data Compression?

There's lots of

New Stuff

From Greenleaf

<http://www.gleaf.com/~gleaf>

visit us today...

Greenleaf Comm++ 3.0

Async communications the C++ way.

Classes for file transfer, terminal emulation, screen drivers, modem and port control, flow control. Supports 16550, smart and non-smart multi-port boards. Supports Windows 95, NT, Windows 3.1x, OS/2 Warp, Novell NASI and 32-bit DOS.

CommLib 5.2 Level 2

Asynch communications for Visual Basic, C, and C++. Supports all popular smart & non-smart multi-port boards, 16550, networked ports. One function call modification to change hardware type. File transfers, modem and port controls, flow controls, terminal emulations, plus hundreds of extras. Supports Windows 95, NT, Windows 3.1x, and 16- and 32-bit DOS. CommLib—the industry standard for 12 years.

No royalties. Lots of examples. Onscreen documentation. Library products come with free source code. 30-day no-questions money-back. Download free test-drive software.

Call today for complete information, demo, or to order. MasterCard, VISA, AmEx and Novus; approved PO's

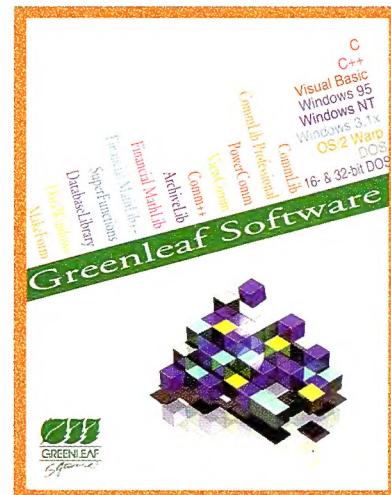
1-800-523-9830

NEW—ViewComm for Windows

Asynch communications debugger for Windows 95 and NT. Peek at RS232 lines, see data and control signals in realtime or review captured data. Flexible trigger system provides start-stop control of display or capture. Graphical display of modem & control signals. Data display in hex, octal, decimal, binary. ASCII & EBCDIC data. Character or hex displays. Searches. Many extras. ViewComm/DOS also available.

ArchiveLib 2.0

Greenleaf proprietary and PKZIP 2.0x compatible compression engines. Visual Basic, Pascal, Delphi, and C/C++ APIs. Flexible archiving including compatibility with PKZIP. Archive to memory or file. Supports Windows 95, NT, Windows 3.1x, OS/2 Warp, and 16- and 32-bit DOS.



Greenleaf Software, Inc.

16479 Dallas Parkway, Suite 570

Dallas, TX 75248

800.523.9830 214.248.2561

FAX 214.248.7830

BBS 214.250.3778 28.8/8/N/1



<http://www.gleaf.com/~gleaf>

email info@gleaf.com

□ Request Reader Service #100 □



HIGH PERFORMANCE

CONVERSION

SCALE-TO-GRAY

IMAGE PROCESSING

THUMBNAILS

PRINTING

SCANNING

FAST DISPLAY

COMPRESSION

and more!

The Best Imaging Toolkit - Is Now Even Better!

AccuSoft Corporation announces a totally new product, **ImageGear**, the next generation in imaging technology. ImageGear, the new version of the AccuSoft Image Format Library™, is more than just an upgrade, it is a **whole new level of the quality and performance** that has made AccuSoft the leading supplier of imaging toolkits. Just for starters, we redesigned all the image filters to add several new features to ensure that our support for image formats is by far the most comprehensive. We even added **file-info executables** for every format we support - and, of course, we added **several new formats**.

ImageGear has a **new and improved API** that makes integration even easier, plus many new features

including totally redesigned display functions for **faster image display**, sub-pixel accuracy, transparency, and an automatic image window with many built-in features.

We also added new **GUI functions** including a comprehensive thumbnail browser, pan and zoom windows, **magnifying glass**, page sorter and more. All with **complete programmability**. We completely revamped the image processing to add **dozens** of new functions and better performance. In addition, we have improved the **scanning**, image compression, display effects, and just about everything else!

So get into **high gear** with ImageGear and remember - the best imaging toolkit makes the best applications - **ImageGear 6.0!**



You can start using
AccuSoft products
today by taking
advantage of our
unique 30 Minute
Delivery Program.

Toll Free: (800) 525-3577
Website: accusoft.com

AccuSoft
High Performance Imaging

Two Westborough Business Park Westborough, MA 01581 Tel (508) 898-2770 FAX (508) 898-9662

©1996 AccuSoft Corporation. All Rights Reserved. All company and brand names are trademarks or registered trademarks of their respective owners. 1. Raster only. 2. License required from Unisys for formats using LZW compression. 3. Optional.

Versions Available

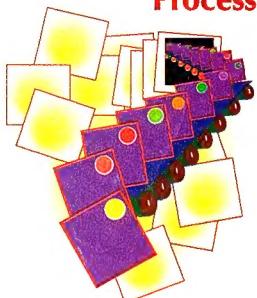
Windows
Win95/NT
MIPS NT
Alpha NT
VBX
OCX
OS/2
SUN OS
Solaris
HP-UX
RS/6000
SGI
SCO
MAC
PowerMac

Over 45 File Formats

JPEG
TIFF
Group III
Group IV
PCX
TGA
PGM
DIB
IMT
DCX
BMP
KFX
RLE
LV
CALS
ATT
CLP
XWD
XBM
CIF
IFF
SUN
PNM
IOCA
GX2
XPM
ASCII
CUT
GEM
BRK
MAC
PSD
MSP
PNG
PCD
IMG
IGF
SGI
ICO
PBM
PPM
MO:DCA
WMF¹
WPG¹
PICT¹
EPS¹
GIF²
ABIC³
JBIG³
DICOM³
and others



STL



Processing Variant Records with STL

19

Linda Kasperek

When you need a map class, you need a map class, and STL provides a good one for many applications.

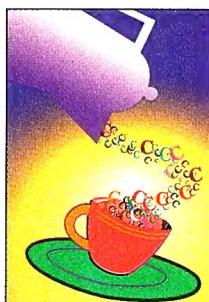
Two STL Books

29

Warren Young

Warren Young reviews two important books for those wanting to learn or become more proficient with STL.

FEATURES



A C++ Chronograph Class

35

Greg Messer

Here's a class for timing program intervals that's powerful in its simplicity.

The Java and C Connection

43

Anil Hemrajani

Sure, Java is a neat new toy. But it's nice to know you can mix in a bit of C code, from time to time, to beef up a Java program.

Floating-point Summation

51

Evan Manning

Addition may be one of the first operations you master when learning mathematics, but it is probably one of the last ones you learn how to do safely on a computer.

In Search of a Portable Screen Library

57

Laura Michaels

If you thought finding a good screen library was difficult, you're right. But author Michaels supplies some useful hints to aid your search.

■ Visit the *C/C++ Users Journal* web site: <http://www.cuj.com> ■

C/C++ Users Journal (ISSN 1075-2838) is published monthly by Miller Freeman, Inc., 600 Harrison Street, San Francisco, CA 94107 (913) 841-1631. Periodicals paid at San Francisco, CA and additional mailing offices. POSTMASTER: Send address changes to *C/C++ USERS JOURNAL*, P.O. Box 52582, Boulder, CO 80322-2582. Subscriptions: Annual renewable subscriptions to *C/C++ Users Journal* are \$34.95 US, \$46 Canada and Mexico, \$65 overseas. Payments must be made in US dollars. Make checks payable to *C/C++ Users Journal*. GST (Canada): L #129065819

COLUMNS

**P. J. Plauger****Standard C/C++: Introduction to <algorithm>**

Programmers should not have to reinvent the wheel, much less test it. STL helps you avoid both, by providing tested implementations of the most commonly needed algorithms.

8

**Dan Saks****C++ Theory and Practice: Abstract Declarators, Part 3**

Sometimes looking at a problem a different way makes it all come into focus. Dan presents an alternative grammar to make elements of C++ syntax much easier to understand.

61

**Bobby Schmidt****The Learning C/C++urve: C->C++ Mutations, Part 4**

In this installment, Bobby wraps up his series on porting C code to C++. Once again he shows that adapting to C++'s stricter rules is painful, but ultimately results in better code.

69

**Pete Becker****Questions & Answers: Little-Known Effects of Defining Constructors**

Find out why constructors and initializers don't always get along, and when to disobey the gods of OOP.

77

**Victor R. Volkman****C/C++ Sources: STL Help on the Web**

You don't have to be a C++ expert to use STL. The World Wide Web is a great place to find tools, tutorials, and moral support.

89

DEPARTMENTS

Editor's Forum	6
Advertiser Index	72
R&D Books	82
Call for Papers	88
New Products	93
We Have Mail	97
Programmer's Market	98

CUJ Online Source Code

Obtain all code published (and some unpublished) in CUJ from:

FTP: ftp.mfi.com in pub/cuj
 BBS: The Courts of Chaos, 501-985-0059;
 EmmaSoft Shareware Board, 607-533-7072;
 Phoenix Chapter ACM Library, 602-821-1162;
 C_BBS (The Netherlands), +31-(0)-4930-23061
 or +31-(0)-4930-20792.
 CompuServe:
 GEnie: GO SDFORUM, Library 7 (R&D Publications)
 In the IBM-PC Roundtable at page 615
 (Keyword: IBMPC).
 Monthly Code Disk: Call Miller Freeman, Inc., Customer Relations,
 1-800-444-4881 or 913-838-7500.

**EDITORIAL**

Senior Editor	P. J. Plauger
Managing Editor	Marc Briand
Contributing Editors	Dan Saks Pete Becker Bobby Schmidt Victor R. Volkman Chuck Allison Martha Masinton Peter Hutchinson
Consulting Editor	Twyla Watson Bogaard
Associate Publisher	Amy Pettle
Publishing Director	Lori White

PRODUCTION

Art Director	Twyla Watson Bogaard
Production Editor	Amy Pettle
Advertising Materials	Lori White
<i>Entire contents Copyright © 1996 Miller Freeman, Inc. No portion of this publication may be reproduced, stored, or transmitted in any form, including computer retrieval, without written permission from the publisher. All Rights Reserved. Quantity reprints of selected articles may be ordered. By-lined articles express the opinion of the author and are not necessarily the opinion of the publisher.</i>	

Printed in the United States of America.

ADVERTISING AND MARKETING

Director of Sales & Marketing	Bill Uhler
Acct. Manager, East/United Kingdom	Ed Day 913-838-7547 eday@mfi.com
Acct. Manager, Midwest	Christine Woodley 913-838-7546 cwoodley@mfi.com
Acct. Manager, West	Julie Thornton 913-838-7541 jthornto@mfi.com
European Advertising Representative	breakout! marketing Duevelsbeker Weg 4 D-24105 Kiel Germany +49 431-801740 FAX +49 431-801797 100332.1704@compuserve.com

Sales Support Manager	Joni Hernly
Circulation Manager	Cherilyn Olmsted

CUSTOMER SERVICE

1-800-365-1364
303-678-0439
FAX 303-661-1885



Advertising: For rate cards or other information on placing advertising in C/C++ Users Journal, contact the advertising department at (913) 841-1631, or write C/C++ Users Journal, 1601 W. 23rd St., Ste. 200, Lawrence, KS 66046-2700.

Customer Service: For subscription orders and address changes, contact C/C++ Users Journal, P.O. Box 52582, Boulder, CO 80322-2582; Telephone 1-800-365-1364 or +1-303-678-0439; FAX +1-303-661-1885; e-mail cujsub@mfi.com.

Back Issues and Article Reprints: 1-800-444-4881 or +1-913-841-1631.

MILLER FREEMAN INC.

Chairman/CEO	Marshall W. Freeman
President/COO	Thomas L. Kemp
Senior VP/ CFO	Warren "Andy" Ambrose
Senior Vice Presidents	David Nussbaum Darrell Denny Donald A. Pazour Wini D. Ragus
Vice President/Software	Regina Starr Ridley
Vice President/Productions	Andrew A. Mickus
Vice President/Circulation	Jerry Okabe

Miller Freeman
A United News & Media publication

From the start, embedded developers have a real time advantage with Phar Lap.

Phar Lap's ETS™ TCP/IP allows your mainframes, workstations, or PCs to communicate with products on the factory floor, in the lab, or at a remote field site.

ETS TCP/IP Software— Harness the Power

The ETS TCP/IP software is part of Phar Lap's TNT Embedded ToolSuite®, Realtime Edition version 8.5 SDK, complete with Realtime ETS Kernel. It is based on a subset of industry standard Win32® and WinSock APIs, the same APIs as Windows® NT® and Windows 95—making it easier and faster to get your product to market!

Ethernet Support— Stable as Ever

Software drivers and sources for:

- 3COM: 3C509
- SMC: WD803, WD8216
- Novell: NE2000 compatibles



For systems with custom interfaces or unsupported cards, the TNT ToolSuite has solutions to make it easy to create your own.

Serial/Dial-up Support— Lots of Horsepower

When designing remote monitoring systems, your system can be connected to a telephone in order to dial up a host computer directly or via an Internet Service Provider (ISP). Either way, it saves you money!

ETS Web Connection— Ahead of the Pack

Phar Lap provides all of the required Web technology including an HTTP server (Web server) and an HTML on-the-fly library for converting raw data into HTML pages for the Web or, if you prefer, just part of your customer's Intranet.

Network Protocols Supported— Putting You on Track

TCP	IP	UDP
ICMP	ARP	RARP
DNS	HTTP	FTP
TELNET	HTML	

As you can see, Phar Lap's TNT Embedded ToolSuite, Realtime Edition version 8.5 SDK has it all: development tools, a Realtime ETS Kernel, and ETS TCP/IP. So stop horsing around and visit us at the "World's Smallest Web Server" URL <http://smallest.pharlap.com> or call a Phar Lap sales representative today!

(617) 661-1510



Embedded Development—Simply on Target™

Looking For Version Control? See What The Experts Have Found.

■ ■ ■ ■ ■ ½ "Source Integrity is more than just a file management system. It performs revision control and simplifies data retrieval, and groups files as projects... so stop reading and get back to writing code. Only this time, take Source Integrity along; it'll make the ride a lot easier." (Thom Duncan—LAN TIMES)

RATING: ★★★★ "One of my favorite features is event triggers... Source Integrity's implementation of this feature is especially great because it lets you build a code block in a simple but powerful scripting language." (SOFTWARE DEVELOPMENT MAGAZINE)

"Would you like to know how SI (Source Integrity) integrates with Microsoft Corp. C++ and Borland International Inc. C++? One word: seamlessly."

(Thom Duncan—LAN TIMES)

"One highlight of the product is a visual merge facility that allows developers to merge different versions with the click of a mouse. 'Sandbox' personal workspaces are supported, where developers can make and test the effects of changes independent of the project as a whole."

(PC TECHNIQUES)

Finding the right version control solution is crucial to maintaining the integrity of your software. You need a proven solution that will guarantee the results you're looking for. Our award-winning software has brought us accolades from respected industry leaders worldwide, establishing *MKS Source Integrity* as the experts choice for version control management.

MKS Source Integrity handles multiple development environments including Borland C++, Borland Delphi, Microsoft Visual Basic, Microsoft Visual C++, PowerBuilder, and Watcom C++. Unlike traditional client/server SCM tools, the MKS solution encompasses distributed development environments, the entire Web team, and intranet developers. If that doesn't convince you, maybe this will: over 450,000 developers worldwide use MKS products to help them accelerate their team's productivity, protect their software assets and guarantee overall source code integrity.



mks
MORTICE KERN SYSTEMS INC.

Mortice Kern Systems Inc.
Tel: 519-884-2251 Fax: 519-884-8861
<http://www.mks.com/solutions/si/cu600/>

Member Netscape

Development
Partners Program.

MKS Germany: TEL +49 711 167140 MKS Nordic: TEL +45 33 25 6555 MKS UK: TEL +44 171 624 0100

MKS is a registered trademark of Mortice Kern Systems Inc. All other trademarks acknowledged.

Managing Change With Integrity

Need a Modern Parsing Tool??



New! Version 2.0

Visual Parse++ is the only tool that applies the techniques of visual programming and RAD (Rapid Application Development) to lexing and parsing technology. Why settle for a text-based tool when you can:

- † See the Syntax Tree
- † See the Parsing Stack
- † See Conflict Trace Trees
- † See Grammar Rule Matches
- † See Errors and Error Recovery
- † See Regular Expression Matches
- † See the Flow of your Input Data
- † Work in one tenth the time!

When your design is finished, write your application using:

C/C++: full support for VC++ 1.5, 2.x, 4.x and BC++ 4.x (16 and 32 bit). Comes with a complete, 100% portable class library (Unix, Mac, OS/2)

Visual Basic: use the powerful parsing VBX/OCX

Delphi: native Delphi support!

Comes with drop-in parsers for HTML, SQL, RTF, Modula 2, C, C++, CONFIG.SYS, and more.

Big claims? We back them up with a 60-day money back guarantee. If it doesn't do what we say, send it back!

Intro price: \$299!

Call (800) 988-9023



950 Shore Crest Road
Carlsbad, CA 92009
Phone: (619) 929-9778
Fax: (619) 929-9848

Request Reader Service #106

Editor's Forum



Okay, folks, it's standards update time again. I'm still jet lagged from several grueling weeks of travel. In late June, I convened the X3J11/WG14 (C) meeting in Amsterdam. Got home for the week of the Fourth, which also included a quick overnight trip to Washington DC. No sooner had I readjusted to Eastern Standard Time but I had to go back to Europe for the X3J16/WG21 meeting (C++) meeting in Stockholm. At the time this schedule was set up, I didn't think it would affect me. Had I But Known, I would have used my limited powers as WG14 Convener to make it more humane.

In any event, the C meeting went smoothly, as it usually does. About 20 attendees, representing six countries, were present. Most are old friends by now and work smoothly together. The committee looks to be on schedule to have the technical particulars of C9X decided by the end of 1996. A formal standard should still occur in time to make that 9X an honest designation.

Without going into a lot of detail, I'll simply say that C9X looks to be the Standard C you've come to know and love with a number of features added primarily for numerical programming. That's largely because the Numerical C Extensions Group was for several years the principal clearing house for extensions to Standard C. The C committee is very conservative, turning down essentially all proposals for which there is no existing practice somewhere. The sad thing, to me, is that any hope of adding class extensions to C, a la C++, is probably dead. The principal champion for this ambitious addition, Bob Jervis at Sun Microsystems, is now caught up in the Java frenzy and can no longer attend meetings.

The C++ meeting went fairly smoothly too, at least through Thursday evening when I left. About 50 attendees, also representing six countries, were present. The goal of this meeting was to complete the changes in response to the first Committee Draft ballot, and to vote out the draft for a second such ballot. Actually, it was the goal of the *previous* meeting to do so — the March meeting in Santa Cruz couldn't quite handle all the issues in time, however.

As far as I know, the committee didn't actually vote out the draft this time. A few issues remain to be handled. More important, several national bodies wanted time to review the updated draft before sending it up the line to ISO. But the committee did commit to a feature freeze. If they hold to their promise to themselves, then Standard C++ may actually be pretty well nailed down by now.

Pete Becker and Dan Saks were also at the meeting, so you should be reading more detail in these pages, from various perspectives, in the coming months. Do stay tuned.

P.J. Plauger

P.J. Plauger
Senior Editor

Announcing a New CUJ Web Service: Our Favorite Source Code Sites

Visit our web page at <http://www.cuj.com>

Ultimate Grid™ v3.0

100% MFC Object Oriented Design



▼ The Premier Grid Control for C++ development now breaks all the barriers. A new 100% MFC Object Oriented design redefines ease of use. Unparalleled power, Unparalleled features, Ultimate performance.

▼ Built-in support for DAO, ODBC, and 3rd party database libraries from Codebase, Faircom and Greenleaf make database integration easier than ever.

▼ Connect directly to any datasource: arrays, memory mapped files, text files, proprietary databases, SQL, real-time, and more. Ultimate Grid gives you total flexibility.

▼ Includes a high-powered memory manager for preloading the grid, and the ability to load data virtually.

▼ Multiple cell types: Bitmaps, checkboxes, radio buttons, drop lists, progress bars, spin buttons, multi-line, auto font size, and more. Too many to list, download our demo for more examples!

▼ Use as a CView, CWnd, CDialog or dialog resource! Fully UNICODE enabled, includes Print/Preview, Find/Replace, and much, much more.

▼ Very simple to use: implement a DAO connected grid with full editing, in just minutes. If you wish, customize the grid's operation. You control everything down to the smallest detail.

▼ Features galore! Inline masked editing, OLE drag and drop, multiple text styles, hints, multiple selection, adjustable column and row height, multi-line cells, multiple data sources, cell joining, multiple fonts per cell, and more!.

▼ Customize the grid to your needs by easily generating your own custom cell styles!

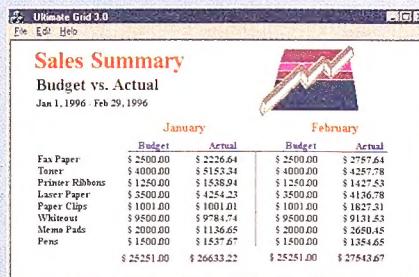
▼ Includes comprehensive source code examples and demonstration code to get you started right away. Includes both a printed manual and comprehensive on-line help.

▼ Royalty-Free!

▼ Site Licensing, license your whole development team for one low price. Call for pricing.

▼ ** Non MFC versions available, call for full product details.

▼ Pricing: \$149.00 / with source \$349.00



Ultimate Grid 3.0		Date
Category	From	Subject
E-Mail	Mark Johnson	Request for product information
	Tammy Anderson	Sales order
	Andrew Peters	What's new?
	Jenny Hols	Please send an order form ASAP
	Randy Duke	Need presentation slides
		Find packager with larger facilities
		Hire more sales staff ASAP
To Do	Sean Winkler	Sales Meeting Next Week
		Increase WWW site bandwidth
Phone	Tony Gardner	Will call back
	Debbie Miller	Need quote
	Bobby J. Townson	Help me, please ASAP
	Jay Sanders	Appointment needed
	Sandra Fisher	Marketing information
		August 1, 1996
		August 2, 1996

Ultimate TCP/IP™

Server and Client TCP/IP Libraries Includes 16 and 32 bit!

▼ A full set of C++ class libraries for developing internet client and server applications. Includes sourcecode for FTP, HTTP, DNS, FINGER, SMTP/POP3 and more!

▼ Build standard or custom applications in just minutes. Modify our sample server and client code to get exactly what you want instantly.

▼ Includes functions and examples to easily create multi-threaded internet server applications that run as true Windows NT services. Our toolkit makes it easy to exploit the latest features of WIN32.

▼ Royalty-Free!

▼ Site Licensing available.

▼ Pricing: \$199.00 / Libraries \$499.00 / Libraries with Source

Ultimate VB Service™

Turn any app into a NT service - Instantly!

▼ Finally, run any application as a Windows NT service. Run Visual Basic, Delphi, MFC, C++, C, Fortran, Pascal, and anything else you might want, even DOS apps!

▼ Well behaved, full featured, and adheres to the NT security model.

▼ Ultimate VB Service allows you to give your projects that extra power, with absolutely no hassle.

▼ Royalty-free!

▼ Pricing: from \$199.00

For a complete list of available products and information

(800) 463-1492

- ▼ Prices shown in U.S. dollars.
- ▼ Call Today! Visa, Amex, MasterCard, or money order.
- ▼ We will deliver via internet, Compuserve, mail or courier.
- ▼ 30 day Money back guarantee. If you do not think these products are the best in their class, we will gladly give you your money back.

202-4800 Dundas Street West, Etobicoke, Ontario M9A 1B1

www.dundas.com

DUNDAS
SOFTWARE

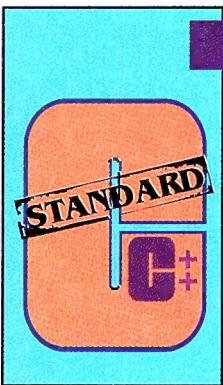
sales@dundas.com

Voice: (416) 239-7472

Fax: (416) 239-2183

Request Reader Service #107 □

240 Portage Rd., Suite 670, Lewiston, New York, USA 14092



P. J. Plauger

Introduction to <algorithm>

Programmers should not have to reinvent the wheel, much less test it. STL helps you avoid both, by providing tested implementations of the most commonly needed algorithms.

Introduction

Last month, I introduced the general topic of algorithms, at least as they are provided by the Standard Template Library, or STL for short. (See "Standard C/C++: Algorithms," *CUJ*, August 1996.) An algorithm in this context is a template function that performs a widely useful function, such as searching a sequence or sorting it. Such functions are typically parametrized in terms of the types of iterators they manipulate to access or generate sequences. (See "Standard C/C++: Iterators," *CUJ*, February 1996.)

STL provides many dozens of algorithm template functions. The vast majority of them are defined in the template <algorithm>, which is by far the largest of the STL headers. (In my implementation of the draft Standard C++ library, it is more than twice as big as the next largest header. But that is partly an artifact of how I chose to chop up the extensive locale machinery.) A handful also appear in the header <numeric>, for some reason.

Even discounting near duplicates, which I discuss below, you can still count over five dozen different algorithms in STL. That's a real treasure trove. Even nicer is the fact that many of these algorithms are hard to get right. It also helps that the template functions automatically adapt when different categories of iterators can profit from different approaches.

P.J. Plauger is senior editor of *C/C++ Users Journal*. He is convener of the ISO C standards committee, WG14, and active on the C++ committee, WG21. He developed *Lib Suite ++*, the Plum-Hall validation suite for the Standard C++ library. His latest books are *The Draft Standard C++ Library, Programming on Purpose* (three volumes), and *Standard C* (with Jim Brodie), all published by Prentice-Hall. You can reach him at pjp@plauger.com.

I don't intend to show you every single algorithm defined by the Standard Template Library, but I do intend to show you many of them. In the process of rewriting the code distributed by Hewlett-Packard, I learned a lot about working with iterators in particular and with templates in general. I think you will find many coding techniques worth emulating here.

This is the first of three installments that describe the algorithms of STL. It is also the most eclectic, I believe. What you see this month is a number of fairly small template functions, each of which does a simple job simply and well. The goal, in each case, is to generate code as good as or better than a good programmer can craft by hand for a given specialization. That's what it takes to make a library that is truly reused, not just nominally reusable.

Non-Sequence Functions

Nearly all the STL algorithms manipulate sequences using iterators. A handful, however, do not. For example, Listing 1 shows two versions of the template function `max`, which returns the larger of its two arguments. Actually, only the first of the two versions is guaranteed to do so. It uses `operator<` to compare the two arguments and determine which is larger.

The second version accepts a *function object* argument as well. As I described last month, such a creature is either a function pointer or an object of some class that overloads `operator()`. In either case, the template function must be able to name the object in a function call expression, with some specified signature. In this particular case, the second form of `max` replaces `x < y` with `pr(x, y)`, where `pr` is the function object.

If you call `max` with the function object `greater<T, T>` (defined in the header

<`functional`>), you get the same behavior as when you omit the third argument. But you can trot in any sort of function object in its place, so long as it accepts a signature of the proper form. Using `less<T, T>` (also defined in <`functional`>), for example, works just fine, but makes a lie out of the name `max`.

Nearly all the algorithm template functions come in two versions. One is hard wired for a widely used predicate, such as `operator<`. The other accepts an additional argument, whose type is a template parameter, so you can supply the predicate in the form of a broad range of function objects. Given the current state of optimization, the duplication of code is often a good idea.

I won't show the other non-sequence template functions, since their implementation is obvious enough given the example in Listing 1. I'll simply mention that they are `min`, with and without a predicate argument, `swap`, and `iter_swap`. Only the last name is less than obvious — it swaps two items given iterators that designate them.

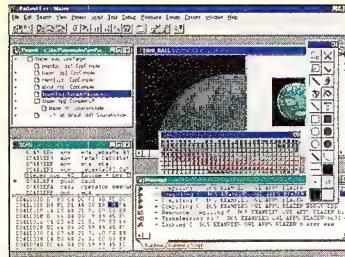
Searching

Listing 2 shows `max_element`, the extension of `max` for a sequence with an arbitrary number of elements. Note that it returns an iterator designating the largest element. In the degenerate case where the sequence is empty (`first == last`), it returns an iterator just past the end of the empty sequence.

Note also that `max_element` comes in two versions, one taking a predicate argument, just like its simpler cousin in Listing 1. From now on, I won't show the more general form. There's quite enough code to look at without the added redundancy. I also won't show the obvious companions, template function `min_element`, with and without a predicate argument.

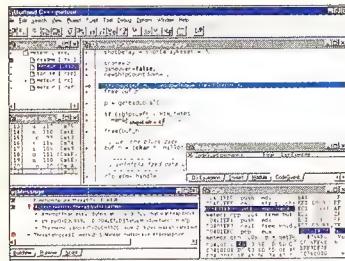
Sometimes you want to consider a sequence as an ordered group of related elements, much like the characters in a string literal such as "abc". Listing 3 shows template function `lexicographical_compare`. It compares two sequences element by element to test whether the first is ordered "before" the second, by the usual rules for ordering words by their spelling. As usual, you can supply a predicate argument to give your own meaning to "before" for two elements (version not shown).

Listing 3 also shows `mismatch`, which compares elements from two sequences until it finds a pair that don't compare equal. It then returns an object of class `pair`



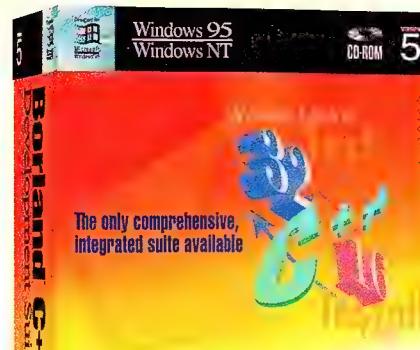
Borland C++ 5.0

Take advantage of the most productive C and C++ environment (included).



CodeGuard

Boost application quality with automatic bug detection and diagnosis (included).



Borland C++ Development Suite

targets Windows 3.1, 95, NT, and DOS



Internet programming with Java

Get the best integrated Java™ development environment including a GUI debugger written in Java. Plus, get the AppAccelerator™ to boost Java application and applet performance by up to ten times (included).*

Cut development time using the New Borland C++ Development Suite 5.0

The only comprehensive, integrated suite available.

The new Borland® C++ Development Suite 5.0 was designed to speed your development. It is the only product to bring together four essential phases of the development process: coding, testing, version control, and install creation—all totally *integrated* for 32-bit Windows development. Plus, get integrated Java tools and AppAccelerator.

Borland C++ 5.0—the most productivity on Windows 95/NT. The suite includes the complete Borland C++ 5.0, designed to bring state-of-the-art innovations in productivity to Windows 95/NT development. Borland C++ 5.0 features a new 32-bit

hosted IDE, parallel 32- and 16-bit development from a single IDE, updated C++ support, updated OWL 5.0 (with more than 150,000 lines of pretested reusable code), ObjectScripting, and MFC library compilation support.

Powerful new ObjectScripting saves you time by automating your build process, integrating tools and utilities, and even adding new features and experts to facilitate compliance with corporate, team, or individual standards. With ObjectScripting, the only limit to productivity is your imagination.

CodeGuard 32/16 automatic bug detection.

Use CodeGuard™ to detect, pinpoint, and diagnose elusive memory and resource bugs in your 32- and 16-bit Windows applications without changing a single line of code. And by double-clicking on your error, CodeGuard will automatically take you to the offending line of code.

Seamless PVCS Version Manager.

Now PVCS version control comes built into your development environment. Check-in code, check-out code, retrieve old versions, visually compare changes in full color, and protect all of your work. Great for both solo and team-based code management.

InstallShield® Express point-and-click install creation.

Use any of

13 prebuilt dialog boxes. Visually build disks, test installation, add uninstall functionality, and create a distribution master.

Find out how much more productive you can be with a suite.

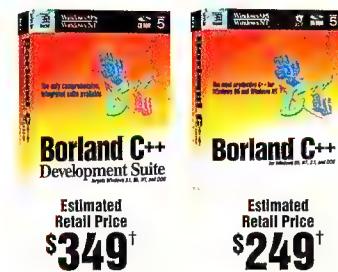
PVCS			
Manage your source code easily and safely with powerful PVCS version control (included).			

InstallShield Express			
Build your installation with point-and-click ease (included).			

Check for page 2-2222

InstallShield Express

Build your installation with point-and-click ease (included).



Borland C++
Development Suite

targets Windows 3.1, 95, and DOS

Estimated
Retail Price
\$349

Borland C++
Development Suite

targets Windows 3.1, 95, and DOS

Estimated
Retail Price
\$249

Get FREE update to Borland C++ 5.0—
now includes MFC.

PLUS, create your own Web site
with free DeltaPoint Quicksite.

Check out Borland Online to find out more
about MFC and new BC++ Development Suite
with Design Tools 5.0.

<http://www.borland.com>

Or for more information call

1-800-336-6464, ext. 50599

CompuServe: GO BORLAND ■ Canada: 1-800-461-3327

Borland

Making Development Easier

Egghead Software
1-800-555-5555

CompUSA
1-800-COMPUSA

Computer City
1-800-THECITY

Programmers' Shop
1-800-421-8006

□ Request Reader Service #108 □

Win 95 without the risk!

System Commander makes it easy to add Windows 95 and as many other OSes to your PC as you want!

Quick and Easy Installation.

System Commander will have your PC ready to add OSes in minutes.

The first reboot brings up a menu of the OSes already installed. Select the one you want and *System Commander* does the rest.

System Commander

Saves You Time and Effort.

As you install new OSes, *System Commander* automatically saves the key files and adds the new OS to the *System Commander* menu.

System Commander makes it easy to evaluate new OSes without giving up the reliability of your existing operating system.



Saves You Money!

Instead of investing \$1000s in new PCs, you can now have as many as 100 OSes on one PC with up to 14 hard drives of any size or type.

System Commander is only \$99.95 and comes with an unconditional 60 day money-back guarantee. And, for a limited time, get FREE overnight shipping when you mention this ad*.

"Highly recommended"
John C. Dvorak 
PC MAGAZINE

Call now!
1-800-648-8266

V Communications, Inc.

 4320 Stevens Creek Blvd., Suite 120-WD
San Jose, CA 95129 408-296-4224
FAX 408-296-4441
www.v-com.com

* If ordered before noon PST. No Saturday delivery. Standard shipping outside US. In CA add \$7.25 sales tax. Offer is subject to change without notice. Product names are trademarks of their companies. VISA/MC/Amex © 1996

□ Request Reader Service #109 □

that stores the two offending iterators. Template function `equal` uses `mismatch` to test whether the two sequences are completely equivalent. Both functions have a version that takes a predicate argument. In this case, however, the predicate replaces `operator==`, not `operator<`.

Another slew of algorithms searches a sequence for various conditions. Listing 4 shows a representative sampler. Template function `find` finds the first element that compares equal to a value you specify. You can replace the value argument with a predicate object, but you must then call the function `find_if`. Template parameter matching is not smart enough always to safely distinguish between a value and a function object.

`adjacent_find` finds two adjacent elements that compare equal, or that satisfy a predicate that you specify with an added argument. Note that this function requires forward iterators. You can't use an input iterator to designate any but the latest, unconsumed element in a sequence.

`count` counts all the elements that compare equal to a value you specify. As with `find`, the version that takes a predicate argument in place of a comparison value is called `count_if`. The return type of `count` has recently been changed to be the type

Listing 1: Template function max

```
// TEMPLATE FUNCTION max
template<class T> inline
const T& max(const T& x, const T& y)
{return (x < y ? y : x); }

// TEMPLATE FUNCTION max WITH PRED
template<class T, class Pred> inline
const T& max(const T& x, const T& y,
             Pred pr)
{return (pr(x, y) ? y : x); }
//End of File
```

Listing 2: Template function max_element

```
// TEMPLATE FUNCTION max_element
template<class FwdIt> inline
FwdIt max_element(FwdIt first,
                   FwdIt last)
{FwdIt x = first;
if (first != last)
    for (; ++first != last; )
        if (*x < *first)
            x = first;
return (x); }

// TEMPLATE FUNCTION max_element WITH PRED
template<class FwdIt, class Pred> inline
FwdIt max_element(FwdIt first,
                   FwdIt last, Pred pr)
{FwdIt x = first;
if (first != last)
    for (; ++first != last; )
        if (pr(*x, *first))
            x = first;
return (x); }
//End of File
```

associated with its iterator parameters. That change requires partial specialization of templates, however, which is currently unavailable in commercial compilers. So I replace the proper return type with `size_t`, which is usually close enough.

`search` searches for a subsequence within a sequence. Note the standard trick of calling yet another template function (`_Search`) to determine the types used to store distances between the different iterator types. Once partial specialization becomes widespread, the use of `_Dist_type` can be discontinued in favor of the notation required for the return type of `count`. As usual, you can replace the equality comparison for pairs of elements by supplying an added predicate argument. `search_n` behaves much the same, except that the subsequence to search for is a repetition of `n` copies of a value you specify.

`find_end` is like `find`, except that it returns the last match in the sequence, not the first. And `find_first_of` is similar to `find`, except that it returns the first element that matches any of the elements in a second sequence. Both functions accept the usual added predicate argument.

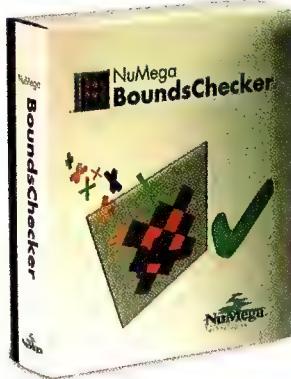
Generators

A handful of template functions perform a specified operation for each element of a sequence. Listing 5 shows these creatures. `for_each` calls the function object you specify for each element of the sequence. Note that it does nothing with any return value. By contrast, `generate` stores the return value of the function object you specify in each element of a sequence, but it lets you specify no argument to the function. `generate_n` does much the same thing for the first `n` elements of a sequence.

`transform` puts the two services together. The unary version applies the function object you specify to each element of one sequence to determine the return value to store in another sequence. The binary version calls the function with two arguments — elements from two sequences — to generate a third sequence.

A simpler way to generate a sequence is to copy an existing one. Listing 5 shows an assortment of such template functions. `copy` copies a specified sequence from beginning to end, while `copy_backward` copies from end to beginning. Knowing which way a `copy` proceeds is invaluable when you have to copy between areas that overlap.

**#1 choice of C++
Windows® developers!**



Find All The Nasty Bugs!

Announcing BoundsChecker™ 4.0

No other automatic error detection tool finds more errors faster!

BoundsChecker 4.0 is the most comprehensive error detection tool available to Windows developers. No matter what level of error detection you need, BoundsChecker finds more bugs faster and more easily in "apples-to-apples" comparisons. Other tools find just the obvious memory leaks and C call errors. With *Smart Debugging*™, BoundsChecker's seamless integration with the Microsoft® Visual C++™ IDE, you find all of these basic errors PLUS all of the other hard-to-find Windows errors.

Only BoundsChecker can:

- **Find more than 85 different types of errors**
- **Validate 5,000 Windows API functions**
- **Validate 70 OLE interface methods**
- **Find OLE interface leaks**
- **Check Microsoft and third-party DLLs and components**
- **Check ActiveX™ Internet controls**
- **Find Windows resource leaks**



**So get the #1 automatic error detection tool for Windows
from NuMega Technologies, the error detection leader.**

Order BoundsChecker 4.0 today... and find all your nastiest bugs!

Call 1-800-4-NUMEGA
(1-800-468-6342)

Or Order Through Our Web Site ...
www.numega.com

9 Townsend West • Nashua, NH 03063
603 889-2386 • 800 4-NUMEGA • Fax 603 889-1135 • info@numega.com • www.numega.com
NuMega Technologies, the NuMega logo, BoundsChecker, the BoundsChecker logo, and Smart Debugging are trademarks of NuMega Technologies.
All other trademarks are the property of their respective owners. Copyright ©1996. All rights reserved.


NuMega
Technologies™

`fill` replicates a value you specify throughout a sequence. `fill_n`, on the other hand, copies the value you specify to the first n elements of a sequence. Note the difference in iterator categories, once again.

Finally, `swap_ranges` exchanges two sequences. It uses the template function `iter_swap` that I showed much earlier. (Strictly speaking, this is not a sequence generator, but sometimes it's hard to characterize a group of loosely related functions.)

Listing 3: Comparison Template Functions

```
// TEMPLATE FUNCTION lexicographical_compare
template<class InIt1, class InIt2> inline
    bool lexicographical_compare(InIt1 first1, InIt1 last1,
        InIt2 first2, InIt2 last2)
    {for (; first1 != last1 && first2 != last2;
        ++first1, ++first2)
        if (*first1 < *first2)
            return (true);
        else if (*first2 < *first1)
            return (false);
    return (first1 == last1 && first2 == last2); }

// TEMPLATE FUNCTION mismatch
template<class InIt1, class InIt2> inline
    pair<InIt1, InIt2> mismatch(InIt1 first, InIt1 last, InIt2 x)
    {for (; first != last && *first == *x; ++first, ++x)
        ;
    return (pair<InIt1, InIt2>(first, x)); }

// TEMPLATE FUNCTION equal
template<class InIt1, class InIt2> inline
    bool equal(InIt1 first, InIt1 last, InIt2 x)
    {return (mismatch(first, last, x).first == last); }
//End of File
```



"the best TSR library I've ever seen."

1994 **EDITORS' CHOICE**

"cream of the crop"

Tom Swan, PC World

"highly useful"

R. Bradley Andrews, Dr. Dobb's Journal

"Every original TSR I converted to CodeRunner went down in memory size, became more compatible and more reliable."

"This software will turn any C programmer into a TSR Michelangelo."

"the size of any program you create with CodeRunner will astound you"

"CodeRunner not only solves problems, it inspires new possibilities. It is destined for the programmer's Hall of Fame."

Joe Campbell, Author of C Programmer's Guide to Serial Communications

"professional development tool that'll let you create compact, fast TSRs"

Gary Entsminger, Micro Comucopia

"excels in its TSR capabilities, coexistence with other DOS applications and support"

Victor R. Volkman, The C Users Journal

CODE RUNNER ®

"the best TSR library I've ever seen."

Stuart Warren, Computer Telephony

"highly useful"

R. Bradley Andrews, Dr. Dobb's Journal

"Every original TSR I converted to CodeRunner went down in memory size, became more compatible and more reliable."

"This software will turn any C programmer into a TSR Michelangelo."

"the size of any program you create with CodeRunner will astound you"

"CodeRunner not only solves problems, it inspires new possibilities. It is destined for the programmer's Hall of Fame."

Joe Campbell, Author of C Programmer's Guide to Serial Communications

"professional development tool that'll let you create compact, fast TSRs"

Gary Entsminger, Micro Comucopia

"excels in its TSR capabilities, coexistence with other DOS applications and support"

Victor R. Volkman, The C Users Journal

- Auto-DISPOSAL of init. code and data
- 300+ funct's in highly optimized ASM
- Hotkeys, schedulers, multiple threads
- Swap/spawn apps, graphics to E/XMS
- Fast background multichannel COMs
- Adv. key stuffing, hi precis. BCD math
- Network, Windows, Desqview friendly
- Quick Start Templates, Online hyp. help
- In depth tech support, "TSR For Hire"
- Trade-in other Libs (limited time offer)

OMEGA POINT, INC.
39 N.Hancock St., Lexington MA 02173


OMEGA POINT, INC.

TEL: (617) 860-0048
FAX: (617) 860-0344

□ Request Reader Service #111 □

Page 12

C/C++ Users Journal — September 1996

Listing 4: Searching Template Functions

```
// TEMPLATE FUNCTION find
template<class InIt, class T> inline
    InIt find(InIt first, InIt last, const T & val)
    {for (; first != last; ++first)
        if (*first == val)
            break;
    return (first); }

// TEMPLATE FUNCTION adjacent_find
template<class FwdIt> inline
    FwdIt adjacent_find(FwdIt first, FwdIt last)
    {for (FwdIt firstb; (firstb = first) != last
        && ++firstb != last; )
        if (*firstb == *first)
            return (firstb);
    return (last); }

// TEMPLATE FUNCTION count
template<class InIt, class T> inline
    // iterator_traits<InIt>::distance_type
    // count(InIt first, InIt last, const T & val)
    // iterator_traits<InIt>::distance_type n = 0;
    size_t count(InIt first, InIt last, const T & val) ///
    {size_t n = 0; ///
    for (; first != last; ++first)
        if (*first == val)
            ++n;
    return (n); }

// TEMPLATE FUNCTION search
template<class FwdIt1, class FwdIt2> inline
    FwdIt1 search(FwdIt1 first1, FwdIt1 last1,
        FwdIt2 first2, FwdIt2 last2)
    {return (_Search(first1, last1, first2, last2,
        _Dist_type(first1), _Dist_type(first2))); }

template<class FwdIt1, class FwdIt2,
    class Diff1, class Diff2> inline
    FwdIt2 _Search(FwdIt1 first1, FwdIt1 last1,
        FwdIt2 first2, FwdIt2 last2, Diff1 *, Diff2 *)
    {Diff1 d1 = 0;
    _Distance(first1, last1, d1);
    Diff2 d2 = 0;
    _Distance(first2, last2, d2);
    for (; d2 <= d1; ++first1, --d1)
        {FwdIt1 x1 = first1;
        for (FwdIt2 x2 = first2; ; ++x1, ++x2)
            if (x2 == last2)
                return (first1);
            else if (!(*x1 == *x2))
                break;
        }
    return (last1); }

// TEMPLATE FUNCTION search_n
template<class FwdIt1, class Diff2, class T> inline
    FwdIt1 search_n(FwdIt1 first1, FwdIt1 last1, Diff2 n,
        const T & val)
    {return (_Search_n(first1, last1, n,
        val, _Dist_type(first1))); }

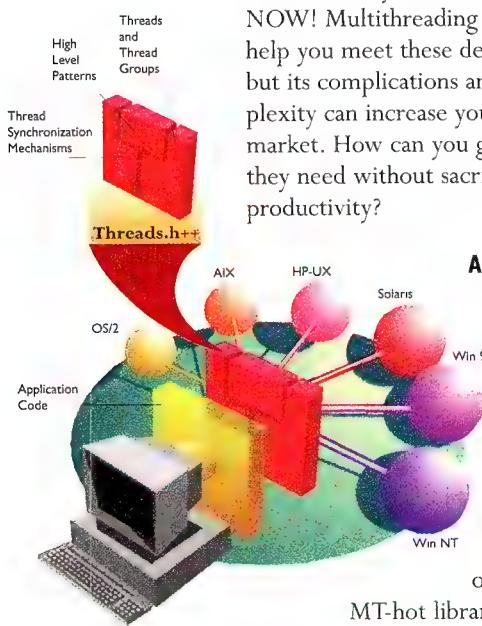
template<class FwdIt1, class Diff2, class T, class Diff1> inline
    FwdIt1 _Search_n(FwdIt1 first1, FwdIt1 last1,
        Diff2 n, const T & val, Diff1 *)
    {Diff1 d1 = 0;
    _Distance(first1, last1, d1);
    for (; n <= d1; ++first1, --d1)
        {FwdIt1 x1 = first1;
        for (Diff2 d2 = n; ; ++x1, --d2)
            if (d2 == 0)
                return (first1);
            else if (!(*x1 == val))
                break;
        }
    return (last1); }

// TEMPLATE FUNCTION find_end
template<class FwdIt1, class FwdIt2> inline
    FwdIt1 find_end(FwdIt1 first1, FwdIt1 last1,
        FwdIt2 first2, FwdIt2 last2)
    {return (_Find_end(first1, last1, first2, last2,
        _Dist_type(first1), _Dist_type(first2))); }

template<class FwdIt1, class FwdIt2,
    class Diff1, class Diff2> inline
    FwdIt2 _Find_end(FwdIt1 first1, FwdIt1 last1,
        FwdIt2 first2, FwdIt2 last2, Diff1 *, Diff2 *)
    
```

Version 1.0

A Simpler Way to Write Portable Multithreaded Apps



Today's sophisticated end users demand more from their applications—power, throughput, responsiveness. And they want it NOW! Multithreading can help you meet these demands, but its complications and complexity can increase your time to market. How can you give them what they need without sacrificing your own productivity?

A WIN/WIN SITUATION
Everybody wins with Threads.h++, the cross-platform solution! Threads.h++ provides portable object-oriented mechanisms to simplify the development of MT-safe and MT-hot libraries and applications. Use the library's encapsulation mechanisms to increase your productivity, while supercharging your apps with improved throughput and responsiveness.

Threads.h++

The C++ class library for writing portable MT-safe and MT-hot libraries

WRITE ONCE, PORT OFTEN
With the Threads.h++ platform-independent, generic multithreading model, write your multithreaded app once, then recompile to run in other environments. Win NT, Win 95, Solaris, HP-UX, AIX, OS/2: the choice is yours!

SYNCHRONIZE YOUR THREADS

You can make your MT-hot app safe when you coordinate interthread synchronization and communication with Threads.h++. The library's complete set of thread synchronization mechanisms includes critical sections; simple, recursive, and FIFO mutexes; semaphores; readers-writer locks; condition variables; and barriers.

FINE-TUNE YOUR APPLICATION

Use the library's thread creation and management classes to take advantage of platform-specific threading features. Check the status of threads, create and destroy threads, and query or adjust the scheduling state of threads to fine-tune your application. Call now to find out how Threads.h++ can increase your productivity.

The Software Parts Company™



Rogue Wave
SOFTWARE

Surf Us
<http://www.roguewave.com>

Call Us
(800) 487-3217
(541) 754-5010

Fax Us
(541) 757-6650

Email Us
threads@roguewave.com

Rogue Wave and .h++, are registered trademarks of Rogue Wave Software, Inc. All other trademarks are the property of Rogue Wave Software or their respective holders.

Sequence Editing

There are various ways to edit a sequence. For the following functions, the code is sufficiently obvious that I won't show it here:

- `replace` replaces each element with the value you specify with another value you specify.
- `replace_if` replaces each element that satisfies the predicate you specify with a value you specify.

Listing 4: continued

```

{Diff1 d1 = 0;
_Distance(first1, last1, d1);
Diff2 d2 = 0;
_Distance(first2, last2, d2);
FwdIt1 ans = last1;
if (0 < d2)
    for (; d2 <= d1; ++first1, --d1)
        {FwdIt1 x1 = first1;
         for (FwdIt2 x2 = first2; ; ++x1)
             if (!(*x1 == *x2))
                 break;
             else if (++x2 == last2)
                 {ans = first1;
                  break; })
        return (ans); }

// TEMPLATE FUNCTION find_first_of
template<class FwdIt1, class FwdIt2> inline
FwdIt1 find_first_of(FwdIt1 first1, FwdIt1 last1,
                     FwdIt2 first2, FwdIt2 last2)
{for (; first1 != last1; ++first1)
    for (FwdIt2 x2 = first2; x2 != last2; ++x2)
        if (*first1 == *x2)
            return (first1);
return (first1);
//End of File

```

Available for Windows 3.1 and Windows 95.

VTOOLSD™

**SIMPLY THE BEST WAY TO WRITE
VIRTUAL DEVICE DRIVERS
FOR MICROSOFT WINDOWS.**

Full sources included for:

- > Plug and Play hardware drivers
- > Network driver templates
- > VCOMM client driver
- > VCOMM port driver outline
- > File system drivers
- > C Run Time Library
- > C++ Class Library
- > Dozens of examples

**GUARANTEED TO SPEED
VxD DEVELOPMENT!**

VTOOLSD is a trademark of Vireo Software, Inc. Microsoft is a registered trademark and Windows is a trademark of Microsoft Corporation.

\$495
FREE TECHNICAL SUPPORT

www.vireo.com

Vireo Software, Inc.
21 Half Moon Hill
Acton, MA 01720
Phone: 508-264-9200
Fax: 508-264-9205
Email: Sales@vireo.com

□ Request Reader Service #113 □

Listing 5: Transforming Template Functions

```

// TEMPLATE FUNCTION for_each
template<class InIt, class Fun> inline
    Fun for_each(InIt first, InIt last, Fun op)
    {for (; first != last; ++first)
        op(*first);
    return (op); }

// TEMPLATE FUNCTION generate
template<class FwdIt, class Fun> inline
    void generate(FwdIt first, FwdIt last, Fun fun)
    {for (; first != last; ++first)
        *first = fun(); }

// TEMPLATE FUNCTION generate_n
template<class OutIt, class Diff, class Fun> inline
    void generate_n(OutIt first, Diff n, Fun fun)
    {for (; 0 < n; --n, ++first)
        *first = fun(); }

// TEMPLATE FUNCTION transform WITH UNARY OP
template<class InIt, class OutIt, class Unop> inline
    OutIt transform(InIt first, InIt last, OutIt x, Unop unop)
    {for (; first != last; ++first, ++x)
        *x = unop(*first);
    return (x); }

// TEMPLATE FUNCTION transform WITH BINARY OP
template<class InIt1, class InIt2, class OutIt, class Bop> inline
    OutIt transform(InIt1 first1, InIt1 last1, InIt2 first2,
                    OutIt x, Bop op)
    {for (; first1 != last1; ++first1, ++first2, ++x)
        *x = op(*first1, *first2);
    return (x); }
//End of File

```

Listing 6: Unique Template Functions

```

// TEMPLATE FUNCTION unique_copy
template<class InIt, class OutIt> inline
    OutIt unique_copy(InIt first, InIt last, OutIt x)
    {return (first == last ? x :
           _Unique_copy(first, last, x,
                       _Iter_cat(first))); }

template<class InIt, class OutIt> inline
    OutIt _Unique_copy(InIt first, InIt last, OutIt x,
                       input_iterator_tag)
    {return (_Unique_copy(first, last, x, _Val_type(first))); }

template<class InIt, class OutIt, class T> inline
    OutIt _Unique_copy(InIt first, InIt last, OutIt x, T *)
    {T val = *first;
     for (*x++ = val; ++first != last; )
         if (!(*val == *first))
             val = *first, *x++ = val;
     return (x); }

template<class FwdIt, class OutIt> inline
    OutIt _Unique_copy(FwdIt first, FwdIt last, OutIt x,
                       forward_iterator_tag)
    {FwdIt firstb = first;
     for (*x++ = *firstb; ++first != last; )
         if (!(*firstb == *first))
             firstb = first, *x++ = *firstb;
     return (x); }

template<class BidIt, class OutIt> inline //!
    OutIt _Unique_copy(BidIt first, BidIt last, OutIt x, //!
                       bidirectional_iterator_tag) //!
    {return (_Unique_copy(first, last, x, //!
                         forward_iterator_tag())); } //!

template<class RanIt, class OutIt> inline //!
    OutIt _Unique_copy(RanIt first, RanIt last, OutIt x, //!
                       random_access_iterator_tag) //!
    {return (_Unique_copy(first, last, x, //!
                         forward_iterator_tag())); } //!

// TEMPLATE FUNCTION unique
template<class FwdIt> inline
    FwdIt unique(FwdIt first, FwdIt last)
    {first = adjacent_find(first, last);
     return (unique_copy(first, last, first)); }
//End of File

```



Watcom C/C++. Performance on any Platform.



Watcom C/C++ accelerates development of lightning-fast, multi-platform 16 and 32 bit applications.

Reliable, high-performance code generation and consistent C and C++ language implementation are delivered across all supported platforms, making it easy to develop applications for several targets from a common source code base. In a single package, Watcom C/C++ provides a comprehensive development environment with the tools, SDKs and libraries you need to create powerful applications for popular PC platforms including Windows®, OS/2® Warp, Windows® 95, and Windows NT®.

*To find out more about
Watcom C/C++, call us at 1-800-265-4555
or visit <http://www.powersoft.com>*

 **Powersoft.**

Watcom C/C++ v10.6 Features

- **Target Platforms:** Windows 95, Windows NT, Windows 3.x, Win32s, OS/2 Warp, OS/2 2.x, Extended DOS, Novell NLM, OS/2 1.x, DOS
- **16 and 32 bit C and C++ compilers with industry-leading optimization technology**
- **Comprehensive MFC support with source, samples, and documentation**
- **Visual Programmer by Blue Sky Software for rapid MFC application development**
- **Multi-platform toolset with a consistent interface across multiple platforms**

- `replace_copy` behaves like `replace`, except that you also specify where to store the modified sequence.
- `replace_copy_if` behaves like `replace_if`, except that you also specify where to store the modified sequence.
- `remove` removes each element with the value you specify, by

Listing 7: Rearranging Template Functions

```
// TEMPLATE FUNCTION reverse
template<class BidIt> inline
void reverse(BidIt first, BidIt last)
{_Reverse(first, last, _Iter_cat(first));}

template<class BidIt> inline
void _Reverse(BidIt first, BidIt last,
bidirectional_iterator_tag)
{for (; first != last && first != --last; ++first)
    iter_swap(first, last);}

template<class RanIt> inline
void _Reverse(RanIt first, RanIt last,
random_access_iterator_tag)
{for (; first < last; ++first)
    iter_swap(first, --last);}

// TEMPLATE FUNCTION reverse_copy
template<class BidIt, class OutIt> inline
OutIt reverse_copy(BidIt first, BidIt last, OutIt x)
{for (; first != last; ++x)
    *x = *--last;
return (x);}

// TEMPLATE FUNCTION rotate
template<class FwdIt> inline
void rotate(FwdIt first, FwdIt mid, FwdIt last)
{if (first != mid && mid != last)
    _Rotate(first, mid, last, _Iter_cat(first));}

template<class FwdIt> inline
void _Rotate(FwdIt first, FwdIt mid, FwdIt last,
forward_iterator_tag)
{for (FwdIt x = mid; ; )
    {iter_swap(first, x);
     if (++first == mid)
        if (++x == last)
            break;
     else
        mid = x;
     else if (++x == last)
        x = mid; }}

template<class BidIt> inline
void _Rotate(BidIt first, BidIt mid, BidIt last,
bidirectional_iterator_tag)
{_reverse(first, mid);
reverse(mid, last);
reverse(first, last);}

template<class RanIt> inline
void _Rotate(RanIt first, RanIt mid, RanIt last,
random_access_iterator_tag)
{_Rotate(first, mid, last, _Dist_type(first),
_val_type(first));}

template<class RanIt, class Diff, class T> inline
void _Rotate(RanIt first, RanIt mid, RanIt last, Diff *, T *)
{Diff d = mid - first;
Diff n = last - first;
for (Diff i = d; i != 0; )
{Diff j = n % i;
n = i, i = j;}
if (n < last - first)
for (; 0 < n; --n)
{RanIt x = first + n;
RanIt y = x;
T val = *x;
RanIt z = y + d == last ? first : y + d;
while (z != x)
{*y = *z;
y = z;
z = d < last - z ? z + d
: first + (d - (last - z)); }
*y = val; } }

// TEMPLATE FUNCTION rotate_copy
template<class FwdIt, class OutIt> inline
OutIt rotate_copy(FwdIt first, FwdIt mid, FwdIt last, OutIt x)
{x = copy(mid, last, x);
return (copy(first, mid, x)); }

// TEMPLATE FUNCTION random_shuffle
template<class RanIt> inline
void random_shuffle(RanIt first, RanIt last)
{if (first != last)
    _Random_shuffle(first, last, _Dist_type(first));}

template<class RanIt, class Diff> inline
void _Random_shuffle(RanIt first, RanIt last, Diff *)
{const int _RBITS = 15;
const int _RMAX = (1U << _RBITS) - 1;
RanIt x = first;
for (Diff d = 1; ++x != last; ++d)
{unsigned long ranm = _RMAX;
unsigned long rann = rand() & _RMAX;
for (; ranm < d && rann != ~0UL;
ranm = ranm << _RBITS | _RMAX;
rann = rann << _RBITS | _RMAX;
iter_swap(x, first + Diff(rann % d)); } }

template<class RanIt, class Ranf> inline
void random_shuffle(RanIt first, RanIt last, Ranf& ranf)
{if (first != last)
    _Random_shuffle(first, last, ranf, _Dist_type(first));}

template<class RanIt, class Ranf, class Diff> inline
void _Random_shuffle(RanIt first, RanIt last,
Ranf& ranf, Diff *)
{RanIt x = first;
for (Diff d = 1; ++x != last; ++d)
    iter_swap(x, first + Diff(ranf(d))); }

// End of File
```

#define NETWORK_DEVELOPMENT

Special Offer:
Request our catalog & receive a Free Squish Ball

1-800-422-6414

800 East Campbell Rd., Suite 199
Richardson, TX 75081

Request Reader Service #115

- remove_copy_if removes each element that satisfies the predicate you specify, by omitting such elements as it copies a sequence.

Listing 6 shows two somewhat more interesting functions that edit a sequence. unique_copy omits all but the first element in any subsequence of elements that compare equal, as it copies the sequence. Unix fans will recognize this as the heart of the classic software tool *uniq*. The code presented here illustrates the use of _Iter_cat to select different approaches based on the category of iterators supplied as function arguments.

Working with input iterators, the function must store an actual copy of an element value. It cannot refer to an earlier value through an iterator once that value has passed by. But working with forward (or stronger) iterators, the function can simply store a copy of the iterator to refer to a value earlier in the sequence. There is less risk that the function will be asked to construct, store, and destroy an object that is expensive in storage or execution time.

In principle, there is no need for the last two overloads. I have found, however, that not all compilers know to select the forward-iterator version for bidirectional or random-access iterators. Thus, the redundant versions are "commented in" as a practical necessity for the nonce.

Template function unique performs in place much the same function as unique_copy. In this case, the latter function can do all the heavy lifting for the former, once unique determines that any rearranging is necessary. It calls adjacent_find for a relatively quick answer to that question.

As usual, both unique_copy and unique have versions that accept an added predicate argument as well.

Rearranging Elements

The last group of algorithms I show this month are lumped together in Listing 7. All of these rearrange the elements of a sequence in various useful ways. reverse and reverse_copy are the simplest. They reverse the elements, either in place or while copying. The former takes advantage of a slight optimization when handed random-access iterators instead of bidirectional ones.

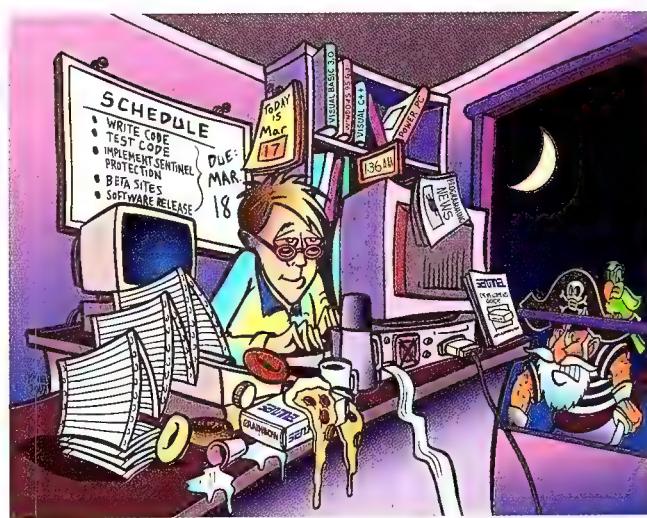
Rotating a sequence is rather harder, at least if you do the job in place. I believe that template function rotate holds the record for taking most ambitious advantage of the ability to fan out on iterator category. I encourage you to study all three versions of this function. They are quite distinct approaches, and each is clever in its own right.

By comparison, rotate_copy is rather tame. In the coming months, you will see how other algorithms take advantage of the greater simplicity of rotating while copying, when the opportunity presents itself to do so.

Finally, random_shuffle does as its name implies. It rearranges the elements of a sequence by performing pairwise swaps of each element with another element chosen randomly. Most of the complexity here stems from a half-hearted attempt to make the algorithm more general. The code endeavors to use the C library function rand to generate a random number big enough to deal with an arbitrary iterator distance type. All it succeeds in doing, however, is to extend a 15-bit function to generate a random unsigned long value. Probably that's good enough for most applications. I haven't yet tried to solve the fully general problem. □

You work too hard to lose sales to piracy

Don't let software pirates steal your revenue



Visit us at COMDEX/CANADA '96
Booth #6149

Sentinel® is the world's leading software protection
Today, software piracy is at an all-time high. It costs developers like you billions in lost revenue each year. That's why more developers worldwide protect software and ensure revenue with Sentinel.

Discover the Sentinel advantage – easy to implement, transparent to your customers, and the most advanced protection. Only Sentinel gives you leading-edge technology, ISO 9002 certified quality and over 99.985% reliability.

Protect your software investment – call for a FREE *Sentinel Software Protection Guide*. Or start protection immediately, order a *Sentinel Developer's Kit*. Each kit has a Sentinel key and the *Sentinel CD-ROM* – it's everything you need to sell more software and make more money.



Call Today!
(800) 852-8569

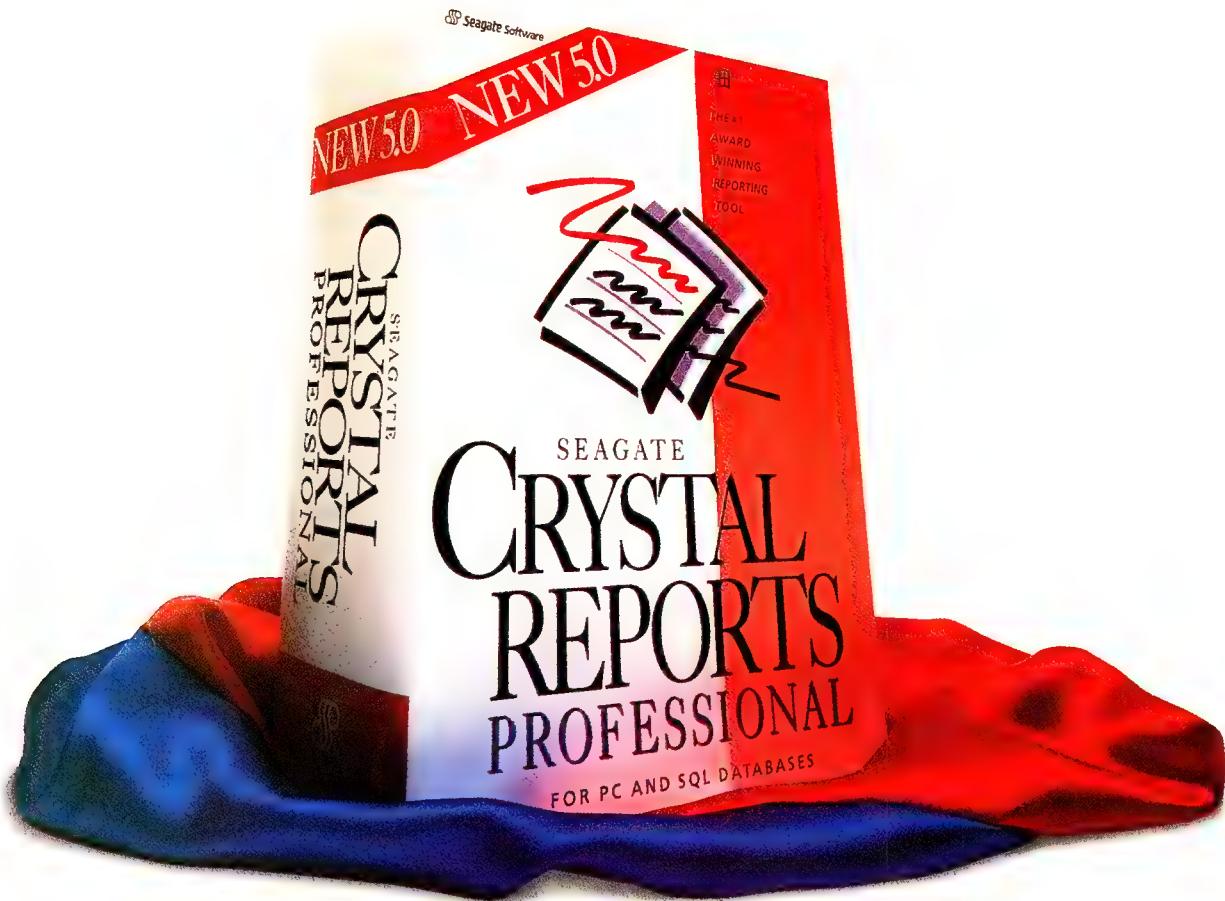
SENTINEL
Software Protection

RAINBOW
TECHNOLOGIES

Visit us on the Internet at: <http://www.rnbo.com> ■ E-Mail: info@rnbo.com
WORLD HEADQUARTERS: 50 Technology Drive, Irvine, CA 92718 ■ Tel: (714) 450-7300 ■ Fax: (714) 450-7450
ASIA/LATIN AMERICA: (714) 450-7300 ■ U.K.: (44) 1932 579200 ■ FRANCE: (33) 1 41 43 2900 ■ GERMANY: (49) 89 32 17 98 0

©1996 Rainbow Technologies, Inc. Sentinel is a registered trademark of Rainbow Technologies, Inc. All other names are property of their respective owners.

□ Request Reader Service #116 □



THE LAST TIME A REPORTER MADE THIS KIND OF TRANSFORMATION, HE WORE TIGHTS AND A BIG "S" ON HIS CHEST.



Get ready for a giant leap in reporting capabilities. Because now there's new Crystal Reports Professional 5.0 — giving you the power to develop and report like never before.

This major release offers an incredible new level of reporting flexibility. With Crystal Reports 5.0, you can create anything from subreports to form-style reports, conditional reports, multiple detail section reports, Web-ready reports, multiple summary cross-tab reports, BackOffice® reports and more.

You also have the added flexibility to publish presentation quality reports directly to the Internet and Intranets using new HTML output options. Plus, you get additional power and control with over 25 new functions, over 15 new ActiveX® (OLE/OCX) properties and many other super new features.

To see Crystal Reports 5.0, visit our Web site: //www.img.seagatesoftware.com/. Then give us a call. You may qualify for the special upgrade price of just \$199 —

that's almost \$200 off the regular retail price. And, you'll receive our 60-day money-back guarantee, which makes your decision to upgrade all the more bullet-proof.

Order your transformed reporter today. Call 1-800-877-2340 dept. CR18.

 **Seagate Software**

Information Management Group
Formerly CRYSTAL A Seagate Software Company

Processing Variant Records with STL

When you need a map class, you need a map class, and STL provides a good one for many applications.

Introduction

Processing variant records, which is often carried out via the brute force approach (a la the big switch statement), can lead to cumbersome code. This code becomes difficult to develop and maintain when the switch statement gets very large. Also, since the switch statement is essentially a "hard-wired" reflection of the record formats to be processed, it provides little flexibility for modifications or additions. Fortunately, C++ gives us a way to handle variant records a bit more gracefully. By storing information about the different record formats and their fields in an STL map, the mechanics can be reduced to a few lines of code. In this article, I show how to use STL to tame the unwieldy switch statement, and how to concentrate data format definitions in one convenient location within the source code.

The Problem with Variant Records

A variant record, as represented in C/C++, is a struct with a union of the variants. The fixed part of the record usually contains a field that specifies the format of the variant part.

A typical variant record structure would look something like this:

```
typedef struct {
    char desc[31];
```

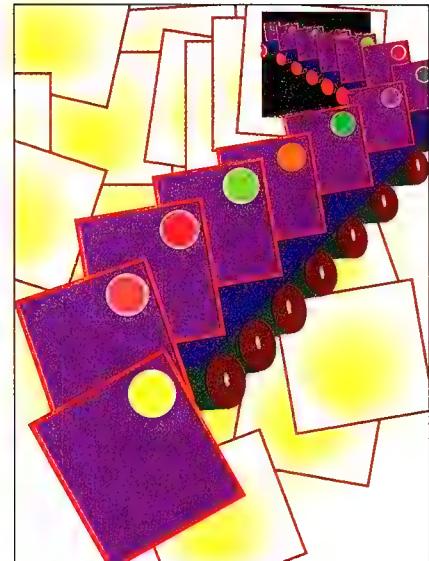
Linda Kasperek has what she considers the ultimate job — developing a Windows-based mass-market software package for a large mid-western financial corporation while working from home. She is pursuing an M.S. in Computer Science at UMass-Lowell, having obtained an M.S. in Statistics in 1993. She has been in software development for 11 years, the most recent three years in MS Windows/C++.

```
char fmt;      // 'A', 'B', or 'C'
union {
    struct {
        // pseudo code
        field1
        field2
        field3
        // to be used if fmt = 'A'
    } A;
    struct {
        // pseudo code
        field1
        field2
        field3
        field4
        // to be used if fmt = 'B'
    } B;
    struct {
        // pseudo code
        field1
        field2
        // to be used if fmt = 'C'
    } C;
} Variant;
} Record;
```

Note that field1 in A, B, and C may not necessarily be the same field.

Variant records in C/C++ are often used to store information from the outside world — commonly a database or configuration file. As an example, I show a simple data file with three different record formats (see Figure 1). In this file, a record spans multiple lines, with one line per field; each record is delimited by the '^' character. Each field is prepended by a field identifier. The file uses these identifiers to allow fields in the variant part to occur in any order, or to be omitted if the variant part is a default value. The first two fields are fixed:

N - name
F - record format



The 'F' prefix in front of the second field indicates that this is a format ID, which identifies the variant portion of the record. Table 1 shows the various field identifiers used within each format.

The structure that will hold these records is shown in Listing 1, as is the brute force approach to reading these records into the structure.

As you can see in Listing 1, the brute force approach requires a pretty ugly switch statement — the code has to branch on each record format and then branch on each field within the record format. This approach would become too unwieldy to maintain if it got any more complicated — and in real life it almost always does! Every time a record format changed, you'd have to muck around with this section of code and you'd probably want to regression test the parts that didn't change. It's easy to introduce new bugs with forgotten break statements, for example. So, you say, there's got to be a better way!

A Solution

There is a more elegant way to handle variant records, using STL's `map` container class and the C/C++ macro `offsetof`. The `map` container provides storage and retrieval of objects via keys. It holds data in the form of Key/Value pairs, where the Value is some object to be stored in the map, and Key is another object used to identify a particular Value object. `offsetof` returns the offset in bytes of a member relative to the base address of its parent structure. The basic idea here is to store the offset as the

Distribute C and C++ APIs—Fast!

EZ-RPC®

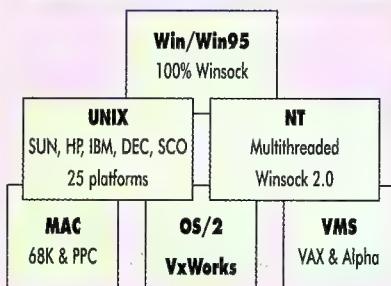
"Client/server made easy"

*Slashes development time
of distributed applications
with "breakthrough"***
third generation
RPC technology.*

* BYTE Magazine
** Patricia Seybold Group



- Distribute database, compute-engine, datafeed, or device APIs
- Rapid prototyping of distributed architectures
- Generates clean readable commented 'C' code
- Supports complex 'C' data structures
- Generates Windows DLLs for VB, PB, and Delphi
- Supports 20+ TCP/IP stacks
- Allows asynchronous and synchronous processing
- Full SPX/IPX for Windows, NT, and NetWare (NLMs)
- 100% ONC/RPC compliant
- No per-user royalties**



NobleNet®
1-800-809-8988

Tel: 508-460-8222
Fax: 508-460-3456
<http://www.noblenet.com>
Email: sales@noblenet.com
337 Turnpike Road,
Southboro, Massachusetts 01772

**FREE
EVAL**

□ Request Reader Service #117 □

Value and to combine the record format and field prefix to make the Key. When the program hands a Key to the map, the map hands back an offset. The offset will be used to point to the corresponding structure

Figure 1: Data file containing variant records

```
NC/C++
F1
$100.99
D04/03/96
^
NUser
F2
D05/17/96
TLa de da!
^
NJournal
F3
$800.25
D12/25/96
THum de dum!
```

Table 1: Field indentifiers

Record Format 1:

\$ – amount
T – description

Record Format 2:

\$ – amount
D – date

Record Format 3:

\$ – amount
D – date
T – description

Listing 1: Variant record struct and switch statement used to populate it

```
typedef struct {
    char Name[21];
    short Format;
    union {
        struct {
            float Amount;
            char Date[9];
        } sFormat1;
        struct {
            char Date[9];
            char Desc[31];
        } sFormat2;
        struct {
            float Amount;
            char Date[9];
            char Desc[30];
        } sFormat3;
    } sVariantRecord;

// switch statement used to process
// above record

switch (recordFormat) {
case FORMAT1:
    switch (prefix) {
        case '$':
            record.sFormat1.Amount =
                atof(p);
            break;
        case 'D':
            strcpy(record.sFormat1.Date, p);
            break;
    }
    break;
case FORMAT2:
    switch (prefix) {
        case 'D':
            strcpy(record.sFormat2.Date, p);
            break;
        case 'T':
            strcpy(record.sFormat2.Desc, p);
            break;
    }
    break;
case FORMAT3:
    switch (prefix) {
        case '$':
            record.sFormat3.Amount =
                atof(p);
            break;
        case 'D':
            strcpy(record.sFormat3.Date, p);
            break;
        case 'T':
            strcpy(record.sFormat3.Desc, p);
            break;
    }
    break;
}
// End of File
```

member into which the data will be copied. This scheme is illustrated in Figure 2.

To instantiate a map container class you need to define a pair of classes — a class to represent keys and a class to represent their corresponding values. You also need to provide for a comparison function object. STL uses the comparison function object, specified by providing a comparison class as the third parameter in the template instantiation, for find and insert operations. The predefined comparison class less<T>, found in STL's function.h, uses operator< to make comparisons. For keys of a built-in type, you get this for free. For user-defined types, there are two alternatives — either write a custom comparison function class, or override operator< in the key and use less<T> where T refers to the key class. In this article, I show how to devise a comparison function class.

Implementing the Variant Map

First, I'll implement the key class. This class is pretty straightforward; it has two attributes, the record format and the field prefix. I also define a copy and default constructor, an assignment operator, a default destructor, and a constructor that takes arguments used to initialize the two attributes. The STL map class requires that the first four functions be explicitly defined. The code for this class (CKey) is

```
break;
case 'D':
    strcpy(record.sFormat1.Date, p);
    break;
}
break;
case FORMAT2:
    switch (prefix) {
        case 'D':
            strcpy(record.sFormat2.Date, p);
            break;
        case 'T':
            strcpy(record.sFormat2.Desc, p);
            break;
    }
    break;
case FORMAT3:
    switch (prefix) {
        case '$':
            record.sFormat3.Amount =
                atof(p);
            break;
        case 'D':
            strcpy(record.sFormat3.Date, p);
            break;
        case 'T':
            strcpy(record.sFormat3.Desc, p);
            break;
    }
    break;
}
// End of File
```

in the `globals.h` header file, shown in Listing 2.

The comparison function class is also simple if you examine the key's attributes separately. The code for this class is also in `globals.h`, just below the implementation of the `CKey` class. This class has just one member: `operator()`. Defining `operator()` enables STL to use this class as a function object. This implementation of `operator()` works as follows: first, it compares the record format attribute for each key and returns TRUE or FALSE, depending on which is the greater. If the values for the record formats are equal, then `operator()` examines the prefix and returns TRUE or FALSE depending on which is the greater.

The value class `CValue` is straightforward as well. It has two attributes, the offset and the type indicator. The type indicator will come in handy when it is time to copy the value into the pointer location determined by the offset. If the value is a string, the program uses `strcpy`. Otherwise, it converts the value into a number using `atoi`, `atof`, or some other appropriate conversion and then copies it

to the target field. This class also meets the prerequisite type requirements by defining the default and copy constructors, default destructor, and assignment operator.

To populate the map, I declare an array of `CKey/CValue` pair objects and then instantiate the map like this:

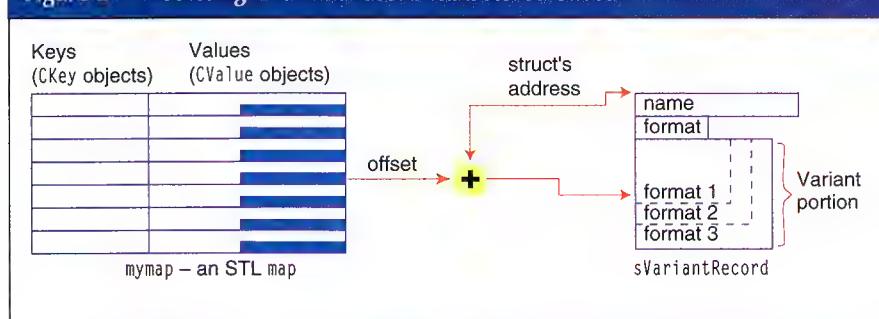
```
mymap m(array, array + 7);
```

`mymap` is a typedef for `map<CKey, CValue, less_CKey>`. The above statement calls a map constructor that takes two pointers to designate an initialization range. This invokes a map initializer that inserts each

pair element in the array (`array`) into the map. Another alternative would be to declare and define each pair and insert it into the map one at a time.

I created the `File` class to encapsulate the reading and managing of variant record files; see `file.h` and `file.cpp` in Listings 3 and 4, respectively. `File` is responsible for getting each complete record from the file. To do this, `File` has to read each line in the record, look up the offset from the map based on the record format and field prefix, copy the value into the corresponding structure member to which the offset points, and return the

Figure 2 Accessing a variant record via a stored offset



There Are Enough Details in the Average Software Project To Make Your Head Spin.

Take Control with Track Record!® The award-winning development tool for tracking bugs, features, releases, and more.

If you're a software developer, project manager, or QA engineer, you need an easy, reliable way to:

- Keep track of what's been done, what's left to do, and where the problems are.
- Stay on top of bugs, features, and changes.
- Figure out who did what, when, and why.

In other words, you need Track Record 2.0, the top-rated development tool from the people who created the revolutionary BRIEF® editor.

Track Record keeps you on top of the details.

Track Record not only tracks bugs from first report through reproduction, fixing, and testing. It also puts you in control of the entire development cycle, helping you improve software quality. It lets you track feature requests...the release a bug fix is in...

bug reports by priority, developer... and virtually anything else you're interested in. It even keeps a running history for future reference.

Most important, Track Record is easy to use. With pre-designed data types, reports, charts, and more, you're ready to take control right out of the box. Working in teams is easier, too. Why? Because Track Record keeps everyone up to date—*automatically*.

What's more, Track Record's fully customizable, object-oriented database gives you the power to track any kind of information. You can even track bugs across several projects at once.

Order now—and find out why the experts rave about Track Record.

Track Record helps all kinds of people take control of the details—and write better software. People like Neil Soane of Mindscape International, who says: "Track Record helps us keep our support people connected and helps us focus on improving our software. We couldn't do without it."

And the trade press definitely agrees: "Track Record is clearly the technology leader—the most flexible and sophisticated of the bug trackers."—Data-Based Advisor, June, 1995

To order, call:

800-343-7308 or 617-267-9743



Just \$195
Risk-free, money-back guarantee
Call for details on client/server version



CodeWizard. The only software available to assist programmers in achieving effective C++ programming — automatically. No need for manual labor; no memorization of applied rules; no crash course in “C++ 101.”

Code Wizard

 Parasoft Corporation



Call now for more information on how to obtain
your copy of *CodeWizard*.

Parasoft Corporation, 2031 S. Myrtle Avenue, Monrovia, CA 91016

Phone - (818) 305-0041, Fax - (818) 305-9048

Website - <http://www.parasoft.com>, Email - info@parasoft.com

populated structure. To look up the offset in the map, File uses the iterator j, declared in mapdef.cpp (Listing 5), and calls the map's find() function with the key j = m.find(key);. find returns the Key/Value pair object. The pair_type template class is a typedef

Listing 2: Defines key and value classes and variant record struct

```
#ifndef _GLOBALS_H
#define _GLOBALS_H

#ifndef _USEMAP
#include <map.h>
#endif

typedef unsigned int BOOL;
const int FALSE = 0;
const int TRUE = 1;

// record formats
const unsigned int FORMAT1 = 1;
const unsigned int FORMAT2 = 2;
const unsigned int FORMAT3 = 3;

#ifndef _USEMAP
// datatypes
const unsigned int FLOAT = 1;
const unsigned int STRING = 2;

```

class CKey {
public:
 CKey() {}
 CKey(unsigned int Format,
 char Prefix) : m_format(Format), m_prefix(Prefix) {}
 ~CKey() {}

 int operator=(const CKey& rhs)
 { m_format = rhs.m_format; m_prefix = rhs.m_prefix; return 1; }
 CKey(const CKey& rhs)
 { m_format = rhs.m_format; m_prefix = rhs.m_prefix; }

 unsigned int m_format;
 char m_prefix;
};

class less_CKey {
public:
 BOOL operator()(const CKey& K1, const CKey& K2) const
 {
 if (K1.m_format > K2.m_format)
 return TRUE;
 else
 if (K1.m_format < K2.m_format)
 return FALSE;
 else
 if (K1.m_prefix > K2.m_prefix)
 return TRUE;
 else
 return FALSE;
 }
};

class CValue {
public:
 CValue() {}
 CValue(long Offset, unsigned int Datatype) : m_offset(Offset),
 m_datatype(Datatype) {}
 ~CValue() {}

 int operator=(const CValue& rhs)
 { m_offset = rhs.m_offset; m_datatype = rhs.m_datatype; return 1; }
 CValue(const CValue& rhs)
 { m_offset = rhs.m_offset; m_datatype = rhs.m_datatype; }

 unsigned int Datatype() { return m_datatype; }

 long m_offset;
 unsigned int m_datatype;
};

typedef map < CKey, CValue, less_CKey > mymap;

for pair<CKey, CValue>, which defines the members first and second, that contain the key and value objects, respectively.

To see all this in action, take a look at the little driver program (Listing 6) I implemented in Visual C++ 4.1. It instantiates a File

```
typedef pair< CKey, CValue > pair_type;
extern mymap m;
extern mymap::iterator j;
#endif // #ifdef _USEMAP

const int cDateLen = 8;
const int cDescLen = 30;

typedef struct {
    char Name[21];
    short Format;
    union {
        struct {
            float Amount;
            char Date[cDateLen + 1];
        } sFormat1;
        struct {
            char Date[cDateLen + 1];
            char Desc[cDescLen + 1];
        } sFormat2;
        struct {
            float Amount;
            char Date[cDateLen + 1];
            char Desc[cDescLen + 1];
        } sFormat3;
    };
} sVariantRecord;

#endif // _GLOBALS_H
// End of File
```

We make software protection easier to handle.



Software Security's new UniKey technology could change the way you think about software protection forever.

Even though our new hardware keys are small, they still offer the most powerful and reliable barrier against software piracy ever developed.

And with competitive pricing, they won't squeeze your margins either.

- Secure protection for multiple applications/versions
- The industry's only patented protection technology
- Protects applications quickly, without programming
- Remote update capabilities
- DOS/Windows/95/NT/Unix/NetWare support
- Time/date-based execution control
- Made in U.S.A.

UniKey makes software protection easier on your customers, your developers, and your budget... but harder than ever on pirates.

Call **1-800-841-1316** for more information, or to order our Developer's Kit, with everything you need to try UniKey protection on your application.

SOFTWARE SECURITY

A new slant on software protection
6 Thorndale Circle, Darien, CT 06820-5421
(203) 656-3000 • Fax: (203) 656-3932
<http://www.softsec.com>



□ Request Reader Service #120 □

object and calls its GetRecord member function until all records have been read. The driver prints out each populated record structure that GetRecord returns. The output appears at the bottom of Listing 6.

Other Uses for this Technique

I've had other occasions to use this scheme — if only I'd been able to get my hands on STL. For instance, I had a task to populate a flat structure with inputs from a

front-end that could only pass it one field at a time — it did not have the capability to pass an entire populated structure. I then used this flat structure to instantiate objects in the back-end for processing. At the time, STL was not available. The field IDs were constant, consecutive integers, so I created an array of structures and could access the array using the field ID as the index. (We avoided using compiler-specific map classes, to maintain portability.) Anyway, having the STL support would have been nice.

Develop our next Wonder.

At Wonderware, we're committed to creating and advancing development in industrial automation Windows-based software. We're the people behind The Factory Suite™ automation software tools. With 50%+ annual growth and an expanding product line, it's no wonder that our clients and our team members are saying "WOW!"

- **SOFTWARE DEVELOPERS**
- **SOFTWARE QA ENGINEERS**

- Solid background in Windows Software, Manufacturing and/or Process Automation • C, C++, MFC, Visual C, Win32 and SDK Toolkit • MS SQL/Server, Delphi V2, VCL and OLE 2.0 a plus

- **MMI/PLC TECH SUPPORT**
- **MANUFACTURING TECH SUPPORT**

- Troubleshooting Windows • PC and PLC programming and/or PLC communication links • Manufacturing position requires supporting implementing manufacturing software
 - Strong knowledge of relational databases • SQL Server
 - Oracle and/or Sybase • NT

The Factory Suite™

Wonderware offers a supportive, small team environment and an excellent compensation package. For more information on these positions, our company and other current opportunities, be sure to check out our website at: <http://www.wonderware.com>

Send your resume to: **Wonderware, Attn: Technical Recruiter/CJ, 100 Technology Dr., Irvine, CA 92718.** As an Equal Opportunity Employer, we value the diversity of the Wonderware workforce.

Wonderware®
Industrial strength software
that's fun to use.

□ Request Reader Service #121 □

Listing 3: The File class

```
#ifndef _FILE_H
#define _FILE_H

#include <fstream.h>
#include "Globals.h"

class File
{
public:
    File();
    ~File();

    void Init(char* FileName);
    BOOL GetRecord(sVariantRecord& record);

private:
    ifstream* m_file;
};

#endif
// End of File
```

Listing 4: Implements file processing

```
#include <stdlib.h>
#include <string.h>
#include "Globals.h"
#include "File.h"

File::File()
: m_file(0)
{
}

File::~File()
{
    if (m_file) {
        m_file->close();
        delete m_file;
    }
}

void File::Init(char* FileName)
{
    m_file = new ifstream;
    m_file->open(FileName, ios::in);
    // int num = m.size();
}

BOOL File::GetRecord(sVariantRecord& record)
{
    char prefix, lpszTokenLine[81];
    short recordFormat;
    float number;
    char* here, *p;
    short datatype;
    CKey key;

    // priming read
    m_file->getline(lpszTokenLine, 81);
    if (m_file->eof()) {
        return FALSE;
    }
    prefix = lpszTokenLine[0];
    p = &lpszTokenLine[1];

    while (prefix != '^') {
        switch (prefix) {
        case 'N':
            strcpy(record.Name, p);
            break;
        case 'F':
            recordFormat = atoi(p);
            record.Format = recordFormat;
            break;
        default:
            key.m_format = recordFormat;
            key.m_prefix = prefix;
            j = m.find(key);
            if (j != m.end()) {
                datatype =
                    (*j).second.m_datatype;
            }
        }
        prefix = lpszTokenLine[0];
        p = &lpszTokenLine[1];
    }
}
```

WWW C++

=Web<ToolKit>TM

- C++ class library
- No HTML experience required
- Supports common HTML 3.0 features
- Supports full HTML 2.0 Spec
- ANSI/ISO compatible
- Easy to Use
- Helps eliminate common mistakes
- Shortens development time
- Operates with all popular versions of ANSI/ISO Standard Template Library and String



CALL 1.800.OBJECT.1

email: sales@objectspace.com

Conclusion

As you may have noticed, using STL didn't get me completely away from the dreaded switch statement. Function File::GetRecord (Listing 4) still needs to do some processing on format types to populate a variant record. The difference between

Listing 4: *continued*

```

switch (datatype) {
    case FLOAT:
        here = (char*)&record + (*j).second.m_offset;
        number = atof(p);
        memcpy(here, &number, sizeof(float));
        break;
    case STRING:
        here = (char*)&record + (*j).second.m_offset;
        strcpy(here, p);
        break;
    default:
        // some error
        break;
    } // switch
}
break;
} // switch
m_file->getline(lpszTokenLine, 81);
prefix = lpszTokenLine[0];
p = &lpszTokenLine[1];
}
return TRUE;
}

//omitted: brute force variant population method. Full source code
//is available on code disk and via ftp, see page 3 for details --mb
//End of File

```

Cmm Brings the Internet to Life!

Everyone wants to write distributed internet applications but no one wants to waste time learning another "revolutionary" language. Cmm was developed with this in mind. Cmm's multi-platform capabilities (Windows 3.1/95/NT, DOS, UNIX, OS/2, and Mac) let you use the same secure scripting language for all your development needs. But unlike those other "hot" internet languages, there is no learning curve with Cmm. If you know C, you know Cmm.

Embed Cmm in your HTML documents, run Cmm CGI scripts on your WEB server, automate browsers with Cmm, or let your users customize your applications with Cmm scripts. Visit our WEB site today and download FREE trial versions of our products. You'll be amazed how fast and easy developing secure internet applications can be.

Cmm Macro Language Toolkit

Let's you incorporate the Cmm language into your application so users can customize your program via your internal calls. Customers will thank you for it.

CEnvi The Standalone **Cmm** Interpreter

Allows you to customize OSs & programs. Make calls to DLLs, kernel APIs, & interrupts. Includes libraries & hundreds of sample Cmm scripts to get you started.

"...CEnvi has virtually replaced ReXX ...If you are a C Programmer, then this product [CEnvi] is really a dream come true."

A. Stevens, Dr. Dobbs Journal

Brian Proffit, OS/2 Magazine

<http://www.Nombas.com/>

Phone: 617 391-6595 Fax: 617 391-3842 BBS: 617 391-3718
nombas@nombas.com 64 Salem Street Medford, MA 02155

□ Request Reader Service #123 □

this switch statement and the "brute force" switch is subtle, but important. The brute force switch requires a separate case clause for every format type, and a separate case under each format for every field type. The improved switch requires a separate case clause only for each field type. This technique basically splits the handling of formats and the handling of field types into two separate, manageable chunks. The format types are more or less

Listing 5: *Creation of Key/Value pairs*

```

#ifndef _USEMAP
#include "Globals.h"

CKey key1(FORMAT1, '$');
CValue value1(offsetof(sVariantRecord, sFormat1.Amount), FLOAT);
pair_type p1 (key1, value1);

CKey key2(FORMAT1, 'D');
CValue value2(offsetof(sVariantRecord, sFormat1.Date), STRING);
pair_type p2 (key2, value2);

CKey key3(FORMAT2, 'D');
CValue value3(offsetof(sVariantRecord, sFormat2.Date), STRING);
pair_type p3 (key3, value3);

CKey key4(FORMAT2, 'T');
CValue value4(offsetof(sVariantRecord, sFormat2.Desc), STRING);
pair_type p4 (key4, value4);

CKey key5(FORMAT3, '$');
CValue value5(offsetof(sVariantRecord, sFormat3.Amount), FLOAT);
pair_type p5 (key5, value5);

CKey key6(FORMAT3, 'D');
CValue value6(offsetof(sVariantRecord, sFormat3.Date), STRING);
pair_type p6 (key6, value6);

CKey key7(FORMAT3, 'T');
CValue value7(offsetof(sVariantRecord, sFormat3.Desc), STRING);
pair_type p7 (key7, value7);

pair_type array [] =
{
    p1,
    p2,
    p3,
    p4,
    p5,
    p6,
    p7,
};

mymap m(array, array + 7);
mymap::iterator j;
#endif //End of File

```

Listing 6: *Test driver program and output*

```

#include <iostream.h>
#include <memory.h>
#include "Globals.h"
#include "file.h"

void main()
{
    File file;
    file.Init("data.dat");
    BOOL bStatus;
    fstream out;
    out.open("out.txt", ios::out);
    cout = out;

    sVariantRecord rec;
    memset(&rec, 0, sizeof(sVariantRecord));

    bStatus = file.GetRecord(rec);
    while (bStatus != NULL) {
        cout << "Name: " << rec.Name << '\n';
    }
}

```

automatically handled by modifying the struct `sVariantRecord`, and perhaps adding a few lines to `mapdef.cpp`. (Adding a format requires adding four lines to `mapdef.cpp`; modifying a format requires no changes to this file.) Field types, which are far less likely to change over the life of an application, are handled in the switch statement.

I suppose I could have pulled this off in C using a couple of associative arrays. But STL let me program intuitively, by defining

Listing 6: continued

```

cout << "Format: " << rec.Format << '\n';
switch (rec.Format) {
  case FORMAT1:
    cout << "Amount: " << rec.sFormat1.Amount << '\n';
    cout << "Date: " << rec.sFormat1.Date << '\n';
    break;
  case FORMAT2:
    cout << "Date: " << rec.sFormat2.Date << '\n';
    cout << "Desc: " << rec.sFormat2.Desc << '\n';
    break;
  case FORMAT3:
    cout << "Amount: " << rec.sFormat3.Amount << '\n';
    cout << "Date: " << rec.sFormat3.Date << '\n';
    cout << "Desc: " << rec.sFormat3.Desc << '\n';
    break;
  default:
    cout << "ho hum" << '\n';
    break;
}
cout << '\n';
memset(&rec, 0, sizeof(sVariantRecord));
bStatus = file.GetRecord(rec);

```

Key/Value pair objects and comparison objects. It was just too portable and elegant a solution to resist. □

Reference

Mark Nelson, *C++ Programmer's Guide to the Standard Template Library*. (IDG Books Worldwide, Inc., 1995).

```

}
out.close();
}

//output:
Name: C/C++
Format: 1
Amount: 100.99
Date: 04/03/96

Name: Users
Format: 2
Date: 05/17/96
Desc: La de da!

Name: Journal
Format: 3
Amount: 800.25
Date: 12/25/96
Desc: Hum de dum!

//End of File

```



LEADTOOLS®
Your comprehensive imaging solution

Scanning
Color Conversion
Displaying
Annotation
Processing
File Formats
Compression
Printing

LEADTOOLS 6 provides 150+ functions for scanning, color conversion, displaying, annotation, processing, file formats, compression and printing. The toolkit has a new **modular design** for building smaller applications.

With LEADTOOLS you can:

- Support TWAIN compliant scanners.
- Convert colors: reducing/expanding color depth, grayscaling, halftoning, creating/merging color separations, and more.
- Support any standard display device with easy to use ultra fast paint, zoom, scroll, crop and pan routines will make your application fly.
- Annotate B&W documents and color images: text, highlight, **sticky notes**, red-line, free hand, blackout and more.
- Process images, using simple or complex regions: spatial transformations, spatial filters, binary filters, histogram modifications, combining of images, despeckling, auto deskewing and more.
- Import/export 40+ of the most popular file formats: BMP, CAL, CCITT FAX, DCX, EPS, GIF*, IMG, ICOA, JPEG, JFIF, JTIF, LEAD CME, MAC, MPT, MSP, PCD, PCT, PNG, PSD, RAS, TGA, TIFF, WINFAX, WMF, WPG, and many more.
- Compress images with multiple options (including CCITT G3 and G4, PNG, JPEG, LEAD's CMP) with the fastest software only JPEG algorithms available.
- Print images, with automatic preprocessing to any supported printer.

LEADTOOLS is offered as an OCX (WIN16 and WIN32), a VBX (WIN16), and a DLL (WIN16, WIN32, and OS/2). The toolkit provides numerous source code examples you can paste directly into your application. LEADTOOLS toolkits come with a 30 day money back guarantee (USA and Canada only) and FREE technical support via phone, fax, BBS, CompuServe and Internet.

Everything You Need To Know About Imaging
<http://www.leadtools.com/>

Download a FREE imaging application built with NEW LEADTOOLS 6!
or call
800-637-1836

LEAD TECHNOLOGIES INCORPORATED
900 Baxter St. Charlotte, NC 28204 704-332-5532 Fax: 704-332-8161 CompuServer: "GO LEADTECH"
LEADTOOLS is available in several versions, not all features are available in all versions.
* License required from Unisys for formats using LZW compression. LEAD and LEADTOOLS are registered trademarks of LEAD Technologies, Inc. All other product names are trademarks of their respective owners.

□ Request Reader Service #124 □

C/C++ Users Journal — September 1996

Make Everyone Happy!



Visit our Web sites:
<http://www.wibu.de>
<http://www.gritech.com>

Order your evaluation kit now: 800-986-6578

WIBU-KEY Copy Protection lets your application work with any of these hardware interfaces with **no modifications** to your application. WIBU-KEY doesn't just do a quick check to see if the dongle is there – it works by encrypting the executable through our custom ASIC. With top notch security and flexibility, along with other great features like remote programming, secure limit counters, and automatic or API-based encryption, you'll quickly see why WIBU-KEY is the best software copy protection available.

We are looking
for Distributors
worldwide



WIBU-SYSTEMS AG
Rueppurrer Strasse 54
D-76137 Karlsruhe, Germany
Phone: +49-721-93172-0
FAX: +49-721-93172-22
email: info@wibu.de, CIS 100142, 1674

GRiffin
TECHNOLOGIES INC.
1617 St. Andrews Drive, Lawrence, KS 66047
Phone: (913) 832-2070, FAX: (913) 832-8787
email: info@gritech.com, CIS 71141, 3624

□ Request Reader Service #125 □

Page 27

Software Developers:

Software Piracy Burns Your Profits.

NSTL Study Rates HASP As Number One!

A recent test conducted by the National Software Testing Labs compared the flagship products of leading software protection vendors. The result? HASP was rated the clear overall winner - and number one in all the major comparison categories. And if the world's leading independent testing lab says HASP is the best, who are we to disagree?

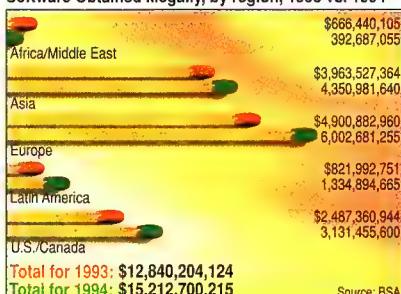
NSTL TEST RESULTS, OCTOBER 1995†

Scoring Category	Aladdin HASP	Rainbow Sentinel	Software Security Activator/M
Security	9.3	6.3	6.2
Ease of Learning	9.1	7.1	7.7
Ease of Use	8.3	7.2	6.3
Versatility/Features	10	8.7	8.6
Compatibility/Power Consumption	6.7	6.5	7.4
Speed of API Calls	0.9	1.2	4.1
Final Score	8.5	6.5	6.6

*For a full copy of the NSTL report, contact your local HASP distributor.

Each year, the illegal use of software consumes nearly 50% of your potential revenues. With the flames of piracy eating away at your profits, can you afford not to protect your software?

Software Obtained Illegally, by region, 1993 vs. 1994



Source: BSA

HASP® is widely acclaimed as the world's most advanced software protection solution. Since 1984, thousands of leading developers have used over two million HASP keys to protect billions of dollars worth of software.

Based on a full-custom ASIC, the new HASP key redefines industry standards of security, reliability and compatibility. And because it's the smallest software protection key in the world, it's also easier on your clients!



Today, more developers are choosing HASP than any other software protection method. To learn why, and to see how easily you can increase your revenues, call now to order your HASP Developer's Kit.

1-800-223-4277

<http://www.aks.com>

ALADDIN

The Professional's Choice

North America Aladdin Software Security Inc.
Tel: (800) 223 4277, 212-564 5678
Fax: 212-564 3377
E-mail: hasp.sales@us.aks.com

Int'l Office Aladdin Knowledge Systems Ltd.
Tel: 972-3-636 2222, Fax: 972-3-537 5796
E-mail: hasp.sales@aks.com

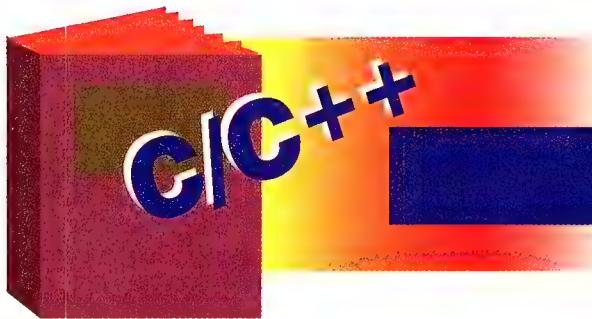
United Kingdom Aladdin Knowledge Systems UK Ltd.
Tel: 01753-622266, Fax: 01753-622262
E-mail: sales@adn.co.uk

Japan Aladdin Japan Co., Ltd.
Tel: +81 426-60 7191, Fax: +81 426-60 7194
E-mail: aladdinj@po.ijinet.or.jp

© Aladdin Knowledge Systems Ltd. 1985-1996 (4/96) HASP® is a registered trademark of Aladdin Knowledge Systems Ltd. All other product names are trademarks of their respective owners. Mac & the Mac OS logo are trademarks of Apple Computer, Inc., used under license. NSTL makes no recommendation or endorsement of any product. The NSTL report was commissioned by Aladdin.



■ Aladdin Benelux 024 641 9777 ■ Aladdin France 1 4085 9885 ■ Aladdin Russia Tel. 095 923 0588 ■ Australia ConLab 3 98985685 ■ Chile Micrologica 2 222 1388 ■ China Shanghai LRI 021 6437 7828 ■ Czech Atas 2 766085 ■ Denmark Berendsen 39 577316 ■ Egypt Zeineldin 2 3604632 ■ Finland iD-Systems 0 870 3520 ■ Germany CSS 201 278804 ■ Greece Unibrain 1 6756320 ■ Hong Kong Hastings 02 5484629 ■ India Solution 11 2148254 ■ Italy Partner Data 2 26147380 ■ Korea DaE-A 2 848 4481 ■ Mexico SiSoft 5 208/472 ■ New Zealand Training 4 5666014 ■ Poland Systemh 61 480273 ■ Portugal Futurimatica 1 4116269 ■ Romania Interact v 64 153112 ■ South Africa D Le Roux 11 886 4704 ■ Spain PC Hardware 3 4493193 ■ Switzerland Opag 61 7169222 ■ Taiwan Teco 2 555 9676 ■ Turkey Mikrobeta 312 467 0635



Two STL Books

reviewed by Warren Young

Introduction

Nothing beats reading a good book for learning a new technology, such as C++'s Standard Template Library (STL). I've read two good books on STL recently, and they both have their selling points. Hopefully, you can afford to get both; but if you can't, this comparative review may help you decide between the two.

Nelson's Book

This book covers the philosophy, use, and intimate details of Hewlett-Packard's (HP) STL implementation. This is a large book, and as with any other large book, one of the first questions I ask is, "Does it really need to be that big?" In this case, I say "Yes." The reason is that, in my view, Nelson's book is really three distinct, complete, and mutually complementary books in one. The first part is actually a long essay (80 pages) which gives an overview of STL. The second is a standard tutorial/reference combination. The third part reproduces the HP STL Specification. In light of STL's almost experimental nature, this last part is closer to being a necessity than just a nice touch. Below, I review each of the three "sub-books" separately.

The book devotes much space to exploring the inner workings of STL, so you really need a good foundation in C/C++ to make sense of it easily. Beginners may still want to read this book, though, for three reasons: 1. STL tutorials are still scarce; 2. Beginners can ignore these explorations and still learn the essentials of STL with this book; and 3. Any beginners reading

this magazine and interested in STL have a good chance of advancing beyond the beginner stage, so the book will eventually become more useful to them.

Book 1: An STL Overview

The first "book" presents an overview of STL, including its history and philosophy.

Title: *C++ Programmer's Guide to the Standard Template Library*
Author: Mark Nelson
Publisher: IDG Books Worldwide, 1995.
Pages: 875, softcover, with disk
Price: \$49.99 US, \$69.99 Canadian
ISBN: 1-56884-314-3

Title: *STL Tutorial and Reference Guide; C++ Programming With the Standard Template Library*
Authors: David R. Musser and Atul Saini
Publisher: Addison-Wesley Publishing Company, Inc., 1996.
Pages: 400, hardcover
Price: \$35.95
ISBN: 0-201-63398-1

Instead of providing the usual quick recap of well-known history, Nelson first thoroughly examines the history that led up to the creation of STL. He starts by discussing the limited genericity possible in C, works his way through C++'s evolving feature set, and ends up at C++'s recent features that make a truly generic and fast library like STL possible.

He then discusses templates. If you already understand C++ templates, you can safely skip this part. For those who can't skip it, rest assured that it's a top-notch discussion of templates as they apply to STL.

Finally, Nelson gives a nickel tour of STL itself, touching on all of STL's major features.

This first book can stand alone to serve as an "STL tasting booth" for those who still haven't decided that STL is something worth learning. If you already know you want to learn STL, this part makes a nice preview for the rest of the book.

Book 2: The STL Tutorial and Reference

The second section is what we might have expected the entire book to be: first a tutorial, then a well-written reference to STL. Indeed, some STL works dedicate the whole book to tutorial and reference. But this second "sub-book" is actually a cut above the rest of the pack. I discuss the tutorial and reference parts in turn.

The tutorial starts with the major sequential container classes: vector, deque, and list, as well as the container adaptors that can make a sequential container look like a stack, queue, or priority queue.

An average book would simply have shown how to use the containers by way of a few examples. What makes this a great book is that after it shows you how to use the container, it shows how HP's STL version really works. Nelson illustrates his explanations with snippets from the HP STL source code as well as clear diagrams.

One of the great things about STL is that it implements each major container in a different way, so that each has its unique performance properties. Accordingly, the next two chapters unveil the very different inner workings of the deque and list containers.

With this detailed treatment the reader can appreciate the underlying mechanisms, which is helpful in making design decisions. It's certainly nice to know that deques perform well when adding elements to both ends, while vectors are only quick for tail insertions. But knowing why these things are so gives you a sure feeling when making design decisions. Nelson puts the topping on this sundae of

Warren Young is a software engineer for Educational Technology Resources, Inc. He has been programming and playing with computers since the mid-80's. He has some web pages about programming (including STL) at <http://www.cyberport.com/~tangent/programming>. You can send him email at tangent@cyberport.com.

GIVE YOUR APPLICATION THE CREDIT IT DESERVES!

"Making a Difference in Transaction Processing"

From Smoky Mountain Technologies... The Multilane Solutions Company

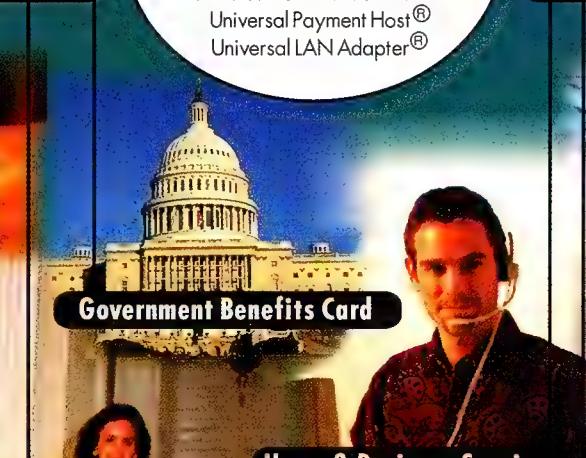
Smoky Mountain Technologies offers a wide variety of Software and Hardware options:

- Universal Payment Software as a stand-alone module or as a 'kernel' to be included with your application.

PRODUCTS



Signature Capture Retrieval



Government Benefits Card



Home & Business Services



Frequent Shopper



Professionals



Mom & Pop Locations



Multi-Lane Retail Locations



AnyCard



Stored Value Card



Credit/Debit/Purchase Card



SMOKY
MOUNTAIN
TECHNOLOGIES, INC.

A UniComp Company

205 Highway 64 West
Murphy, North Carolina 28906

- DOS
- Windows
- Unix

We Provide Payment Solutions for:

- Software Developers
- Cash Register Dealers
- POS Payment Providers
- Anyone Selling into the Multilane Market

Easy Interface • Developer's Kit • Unlimited Phone Support • Installed in Thousands of Locations • Currently Used by Many Developers

□ Request Reader Service #127 □

Please call for more information: (704)837-6079 or Email: info@smokymtn.com

sequential container information by pitting all three containers against one another. This "showdown" of containers demonstrates the space and time tradeoffs clearly.

The next chapter briefly covers STL's container adaptors. Adaptors give a container an interface more appropriate to the task at hand. One common use is turning a vector into a stack, for example.

The next chapter covers STL's four associative containers. As in previous chapters, Nelson covers the implementation underlying the containers, which in this case is a red-black tree. While the explanation is useful, I wonder if the book wouldn't be improved by covering the interface and use of the red-black tree class itself. It seems to me that users may find themselves wanting to derive their own classes from STL's version instead of writing their own.

The next two chapters in the second section cover allocators and iterators, which provide uniform memory management and access to the containers. These chapters cover the subject well, but I can't say it's one that is all that exciting to me. These chapters aren't boring, but aside from being as complete as the rest of the book, they aren't otherwise noteworthy. Maybe those wanting to do things like database and special memory access through custom allocators and iterators will find this chapter more indispensable.

The next chapter gives a very brief overview of STL's algorithms. It basically just tells you what's available so you can find more information in the reference section. Last but not least, Nelson discusses function objects. These clever object-oriented replacements for function pointers are key to many portions of STL.

The final chapters of the second section are a reference for STL, and as far as I've been able to tell, cover everything in HP STL. In addition to expanding upon the tutorial's topics, the reference covers sorting and searching, as well as set, heap, sequence, basic numeric, and miscellaneous operations.

Book 3: The STL Specification

This final "book" is a reprint of the original HP STL specification. This section is also quite valuable; it gives you an "STL bible" to refer to when your questions and problems outstrip the rest of the book's scope.

Criticisms and Kudos

Although I really like the book, I should point out a few of the book's faults and tradeoffs. The first tradeoff to keep in mind is that (probably due to timing issues) the book covers HP STL, which the pages of this magazine have shown to be very different from the emerging draft Standard STL. A serious fault is that the text rarely refers to the STL header files, so it's harder than it should be to find out which headers to include to get particular STL facilities.

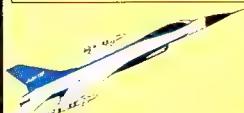
The book more than makes up for these problems, however, with some really nice touches that make it truly great. For example, Nelson lists what each of the containers require of their contained objects. Documentation of this sort is dear to me, as I've frequently been forced to do without it.

As another nice touch, Nelson hints at many possibilities for extending STL through objects such as custom iterators and allocators that use something other than the heap for storage. If you wonder what this non-heap storage might be,



SERIAL I/O

The ONLY Serial Communication Libraries, DLLs, & Tools You Will Ever Need For Windows 3.x, Windows 95, Windows NT, and MS-DOS. Accept No Less Than The Best.



COMM-DRV/LIB™
Professional Serial I/O Libraries & DLLs for Windows 3.x, Windows 95, NT, & MSDOS

- Supports ALL languages, tools, or applications that can call the Windows API. Very Easy To Use!!
- Complete source code to all libraries & DLLs.
- Same API for Windows 3x, Windows 95, NT, & MS-DOS.
- High level Hayes compatible modem functions.
- X, Y, Zmodem file transfers on multiple ports concurrently.
- Asynchronous & timed callback to user functions on different serial communication events(modem signals, buffer counts, character reception, and more).
- Supports Visual C/C++, ANSI C/C++, QuickBasic, Visual Basic (MSDOS & Windows), Assembly, Access, etc.
- Transmit data from user callback on any event(important for multidrop applications).
- Extensive scanning of input character stream.

WCSC Price List

COMM-DRV/Lib(DOS,Win3x,Win95,NT)	\$189.95
COMM-DRV/Lib(Competitive Upgrade)	\$ 99.95
COMM-DRV/VxD(Introductory Offer)	\$ 99.95
COMM-DRV/Dos	\$ 99.95
COMM-DRV/Win	\$ 99.95
4Port Card(16450)	\$110.00
4Port Card(16550)	\$170.00
8Port Multiport Card(16550)	\$225.00
Software Combo(Lib,VxD, Dos, Win)	\$399.95
Complete Combo(Sftwe+4Port/16550)	\$499.95
BBS SYSOP Combo(Dos+4Port/16550)	\$199.95

Features Supported By All Products

- Unlimited number of ports active concurrently.
- Hardware/Software Flow Control.
- 8250/16450/16550 Auto detection/Up to 115200 baud.
- Product Customization.
- 100% Port re-enter code(Important for multitasking).
- Remap/Change baud rates/divisor.
- Support most intelligent/all dumb multiport cards(Arnet, Digiboard, Boca, GTEK , AST, & more).

COMM-DRV/VxD™
Ultra High Speed Serial I/O VxD For Windows

- Windows DLLs with API calls into the 32bit VxD.
- MS-DOS high speed libraries for calling VxDs from a DOS box.
- Handles all serial communication interrupts at ring 0 in 32 bit mode.
- Automatic detection of 16550 UARTs/Uses transmit & receive FIFOs.
- May be called from Windows & MS-DOS concurrently.
- Baud rates in excess of 115.2K underWindows!!!
- Integrates seamlessly with all COMM-DRV products.

COMM-DRV/DOS™
MS-DOS Serial I/O TSRs & Utilities

- True DOS device driver that allows serial ports to be opened like files under MS-DOS & Windows.
- Provides FOSSIL interface for all supported cards(Supports all BBS, Mailers, etc [PCBOARD, Wildcat, FrontDoor, DoorWay, Renegade, etc]).
- Provides Int14h, Int21h, & DAM interfaces.
- Compatible with Desqview and Windows.
- Real time serial port monitoring to screen and disk.
- Mini-BBS for file transfer/Spawnable file transfer engine.

COMM-DRV/WIN™
Windows Replacement Serial Communication Driver

- True Windows communication device replacement. Use your favorite Windows communication applications on all multiport cards we support.
- Allow the use of more than 9 serial ports under Windows by allowing users to remap any serial port on the fly.



Visa/Mastercard/Discover/AMEX/Checks/Approved P.O. **WWW:** <http://www.wcscnet.com/home.htm>

□ Request Reader Service #128 □

C/C++ Users Journal — September 1996

Page 31

Nelson provides some examples: database access and "special" memory access like MS-DOS's EMS memory.

Nelson minimizes the book's bulk through a form of reuse: he uses a single code example to illustrate a series of points in turn.

Finally, the chapter on associative containers attempts to shed some light on why STL's creators chose to implement associative containers as red-black trees instead of hash-based structures. His argument may not end the debate, but he does illuminate the issues well.

Other Considerations

As its back cover indicates, the book focuses on PC compilers, but readers should find it easy to identify and compensate for its few minor platform dependencies. For example, the code samples each have a commented `#define` at the top to make them compile with Borland C++ 4.x. These are easily removed. Another PCism: when the book displays pointers, they are clearly of the 16-bit, "short" variety. Beyond these sorts of things, I believe the code should be portable to any STL-compatible compiler (with the caveat that the diskette is MS-DOS-formatted.)

Musser and Saini's Book

As I hinted in the previous review, I expect a book of this kind to consist of two main parts, a tutorial in the basic-to-intermediate material first, and a complete reference second. Musser and Saini follow this well-established formula.

As you might expect from its (hard) cover, this book takes a more formal approach than Nelson's. It's certainly not from the ivory

tower, but this book has the standard properties of a formal work: no light conversational humor in headings and text, lots of Big-O notation, and platform independent discussions. If it had a lighter writing style, it might be easier to read, but more formal writing has its advantages, too. For example, it's easier to skim the book by glancing at headings, because you don't have to mentally translate their meanings.

The book does not attempt to teach the use of templates, so readers will need to know how to use others' templates before tackling this book.

A note on accuracy: authors are only human, so they make mistakes. Addison-Wesley recognizes this and thus makes its errata sheets easily available. The errata sheet for this book is available on the Web at <http://www.cs.rpi.edu/~musser/stl-book-errata.html>.

The Tutorial

Musser and Saini get right to their subject. Unlike Nelson, they don't try to "sell" you on STL first; if you bought the book, they assume that you really do want to learn STL. They also don't bother recounting the converging histories of STL and C++. I give Musser and Saini points for a better overview of STL, though — they provide less trivial examples and more of them.

The tutorial is one place where the book's theoretical outlook becomes obvious: the first tutorial chapter is on iterators, the glue that binds everything else together. Notwithstanding my lack of excitement about iterators (see previous review) I think this may be a better strategy than Nelson's. Without a good understanding of iterators, fitting STL algorithms and containers together can be quite frustrating.

The next chapter is on STL's algorithms. It might seem more prudent to cover containers next — after all, how can you use algorithms without containers? But you can illustrate a lot of algorithms with simple vectors, which is what the authors do. Also, the basics of vectors were covered in the introductory material. The chapter actually covers all of STL's algorithms well, and with each algorithm (or group of closely related algorithms) it includes a well-chosen example program as illustration.

Next come STL's sequential containers, all in one chapter. This arrangement makes sense because the three sequential containers belong to a "family of abstractions," which lets the authors describe only once a feature common to all of the containers. The next chapter covers associative containers in much the same manner.

The next four chapters quickly cover function objects and the various adaptor types. It is interesting to note that this book considers the adaptors separately from the facilities they adapt, whereas Nelson deals with them side-by-side.

The Example Programs

Again, Musser and Saini outdo Nelson in the examples category: each of the next four chapters is built around a largish example (as compared to the rest of the text's examples) that solves a problem by using STL. We usually like to read an example that does a lot for its size, and the authors succeed handily in creating such an example here. All of the programs are also based on a common theme, working with anagrams; as a result, they all hang together well, and between programs improvements follow a natural progression.

The next chapter focuses on defining an iterator class. This is a nice, mildly evangelistic touch on the authors' part — evangelistic because they're (naturally) trying to prove the extensibility of STL.

Get programming advice from the experts — in seconds!

Over 30 of your favorite
programming magazines
on one CD-ROM!



AI Expert
BYTE
C/C++ Users Journal
CLIENT/SERVER COMPUTING
Data Communications
Database Programming
& Design
DBMS
Delphi Informant
Dr. Dobb's Journal
Embedded Systems
Programming
EXE
IBM Systems Journal
IEEE Computer
IEEE Computer Graphics &
Applications
IEEE Expert
IEEE Software
LAN Magazine
LAN Times
Microsoft Systems Journal
NetWare Solutions
Network Administrator
Network VAR
Open Computing
OpenStep Journal
Oracle Magazine
OS/2 Magazine
Paradox Informant
PC AI
Software Development
Software Magazine
Sys Admin
UNIX REVIEW
Windows Developer's Journal

Here's what you get...

- ◆ A 3-year collection of articles, product reviews, design tips, algorithms, short-cuts, and more — the complete text — from over 30 leading programming publications.
- ◆ 1,000,000 + lines of published source code
- ◆ Time-saving search software
- ◆ 8 prominent software development guides and books — **FREE!**
- ◆ Money-back satisfaction guarantee

Only \$149

Call 1-800-370-6717
Outside the U.S. call 914-968-7008
<http://www.i-mode.com>

□ Request Reader Service #129 □

They then take a quick step out of the classroom and into the "real world." Here the authors discuss some of the how-to issues that come up when trying to get STL to do some nonobvious things. One of the examples show how to reduce the template code bloat problem with a novel application of polymorphism.

The Reference

The book's reference section is, like the rest of the book, quite "standard." It is divided into chapters roughly parallel to the tutorial's, and each subsection discusses the details of a particular facility. I didn't find it lacking in any way. One thing covered in this section that isn't covered in the tutorial are allocators. The reference provides enough information to let you write your own allocators if the necessity presents itself.

Criticisms and Kudos, Part 2

I do have a criticism of the book. I'm really unimpressed with the illustrations in this book. They're not inaccurate, but they just don't communicate their information as well as Nelson's do. Fortunately, I think this book doesn't need illustrations as much as Nelson's book does, so it doesn't suffer all that much. I still think that they could have been clearer, though.

Now here are some nice touches: though the rest of the book uses the HP reference implementation of STL in its examples, the introduction to the book does explain the header changes that the ISO C++ committee made. These particular changes will probably cause some headaches, and this feature should help alleviate the suffering.

The tutorial chapter on iterators offers a table of containers paired with the iterators they return. This makes matching up containers with compatible algorithms easier.

The chapter on associative containers doesn't spend much time arguing for red-black trees over hash tables, but the book does mention that there are STL-based hash containers available from the main STL site (<ftp://butler.hpl.hp.com/stl>), which was a nice revelation. The other sites mentioned in the same appendix are also worth a visit, and their inclusion is yet another nice touch.

Comparing the Two Books

These two books differ in philosophy, in a couple of fundamental ways. For one thing, Nelson's book tries to do more than Musser and Saini's does. Besides the reprint of the STL spec, the main extra Nelson offers is a running exploration into the workings and design of an implementation of STL. However, in doing so, he constrains the book somewhat to one implementation and one platform.

On the other hand, Nelson's coverage gives us an insight into STL that, once obtained, is hard to imagine doing without. This added insight gives the designer in all of us information to help us choose between similar STL facilities, as well as some good lessons in design. Plus, it's really easy to step past the dependencies if they get in your way. In all likelihood, future implementations of STL will look much like the HP reference implementation, and the PC dependencies are easy to spot and fix.

It all adds up to this: I think Nelson's book is a great introduction to the library, and that it provides an insight into some of the "Zen" of STL, which Musser and Saini's book does not. Also, Nelson's inclusion of the STL spec and a decent reference makes his book useful beyond the initial learning period.

Does that mean the other book is useless? No, not by any means. It's just *different*. For one thing, I think that Musser and Saini have

done a better job of organizing their book. The organization of Nelson's book is okay, but I get the impression it wouldn't be as easy to use as a reference. By better organization I mean that it is easier to predict where to find coverage of a certain feature in Musser and Saini. In particular, Musser and Saini make a clear separation between the straight reference material and the tutorial material. Of course, Nelson's organization has its points as well: his reference section contains examples as well as the tutorial, which is nice when you're still learning how to use something and need both reference and example material at the same time.

I mentioned that Nelson's book covers a working implementation of STL. I personally consider this a plus, but perhaps not everyone will. If you don't have the time or inclination to understand the details of an STL implementation, you may be better off with Musser and Saini.

The Verdict

I think Nelson's book makes a better tutorial, and that Musser and Saini have the better reference.

If you're a non-PC user who can't stand PCisms, stay away from the Nelson book. If you're more interested in the theory of STL, start with Musser and Saini's book.

I'd really hate to choose between the two books. But if you must, I'd recommend doing what I did, if you can: learn with Nelson, and if you decide you need better reference or how-to material, get the Musser and Saini book. □



Quadron development tools help the world communicate

Use Quadron's rapid development software with IBM ARTIC co-processor communication cards to offload PC protocol processing. You'll gain multiple ports, high data rates, & multiple protocols on a single card.

X.25, LAP-B, HDLC/SDLC, Bisync, Async and your own custom protocols.

Worldwide, Quadron's software is serving people with applications like process control, gaming, automatic teller systems, stock market data feeds, telephony, airline reservations, subway ticketing, emergency dispatch & robotic control.

How can we help you?

<http://www.rain.org/~quadron/Q.html>



Quadron®

209 East Victoria Street
Santa Barbara, CA 93101
telephone 805-966-6424
fax 805-966-7630



□ Request Reader Service #130 □

Instant Access to Six Years of *C/C++ Users Journal*

Stop fumbling through your magazine heap

The CD's easy-to-use interface gives you powerful, full-text search capability. Find the articles you need instantly.

Never buy another back issue

It's all here: whether you began subscribing last month, last year, or at the beginning of the nineties, you'll have six years of *C/C++ Users Journal* at your command.

Don't reinvent the wheel

C/C++ Users Journal articles supply code that you can drop right into your current application. From tiny functions to class definitions, fellow programmers share the fruits of their labor with you.

The *C/C++ Users Journal CD-ROM* contains six years of the leading C/C++ programming magazine. Hundreds of articles, most with source code, provide practical technical information on debugging, algorithms, class libraries, image processing, portability, optimization, real-time and embedded systems, ANSI C and C++ standards — and much more.



Order Today!

(800) 444-4881

Call 24 hours a day — 7 days a week
Or complete the form and return to :

FAX: (913) 841-2624

E-MAIL: orders@mfi.com

INT'L: Use mail, fax, e-mail,
or call (913) 841-1631

FULL MONEY BACK GUARANTEE. If you don't find the *C/C++ Users Journal CD-ROM* worth the price, we'll return your money — no questions asked.

Please send me a *C/C++ Users Journal CD-ROM* today! My price is \$49.95 (plus shipping: \$2.00 U.S./Canada; all other countries \$10.00) Sales tax added in the following states: CA (8.5%), GA (6%), IL (6.25%), KS (6.9%), NY (8.25%), TX (8.25%).

Name _____

Address _____

City _____

State _____

ZIP _____

Country _____

Phone _____

Check

VISA

Mastercard

Amex

Credit Card# _____

Exp. Date _____

Signature _____

C/C++ Users Journal CD-ROM

1601 West 23rd Street, Suite 200, Lawrence, KS 66046-2700

A C++ Chronograph Class

Here's a class for timing program intervals that's powerful in its simplicity.

A Programming Problem Solved

One coding task I find myself doing over and over in my programs is timing. I time the program as a whole, I time the individual phases of the program, I time the I/O throughput, and I estimate the program's completion time. That timing code is just complicated enough to make me think carefully while writing it — and just tedious enough to make me want to never write it again. What I need is an easily reusable, plug-in tool to simplify those timing tasks for me.

To fill that need I have created a C++ Chronograph class, a software stopwatch to handle my repetitive timing requirements. The Chronograph class provides the same functionality as a good digital stopwatch. It will give the elapsed time, split times, and lap times. You can start, stop, restart, and reset the Chronograph. It can break out the hours, minutes, and seconds for you. Best of all, several Chronograph's can run at the same time without slowing down your program one bit.

Chronograph Functionality

The Chronograph can be in one of two modes, running or stopped. Chronograph::start does one of two things, depending on whether the

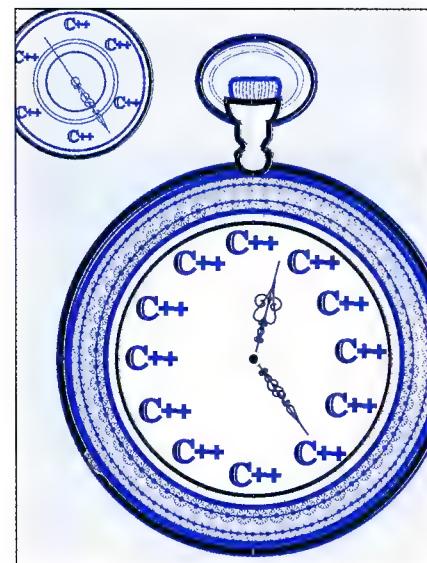
Chronograph is running or stopped. When running, Chronograph::start will restart the Chronograph at zero and leave it running. When stopped, it restarts the Chronograph from zero. I coded Chronograph::start to work this way because that is how a stopwatch works. The return value of this member function is the elapsed time.

Calling Chronograph::elapsed or Chronograph::split gets the elapsed time of execution up to the current moment. The elapsed time is the amount of time since the Chronograph was started. Splits are the same as the elapsed time, so function split just calls elapsed.

Chronograph::elapsedHMS breaks down the elapsed time into hours, minutes, and seconds. Using this function allows you to report "5 hours, 32 minutes, 10.25 seconds" instead of "19930.25 seconds".

Laps are different from the simpler elapsed time. If you call Chronograph::lap it will return the amount of time since Chronograph::lap was last called. Laps are the periodic intervals that exist within a program. Like a runner doing laps on a track, programs loop and perform other operations that have distinct starting and stopping moments. For example, you can keep track of the time spent processing 1,000 database records, so you can estimate the program's completion time, or keep track of the time spent in each phase of a multi-phase process.

Chronograph::stop puts the Chronograph in the stopped mode. The Chronograph stores the clock value when it was stopped and uses it to adjust the start time at the next restart. The amount of time spent in the stopped mode does not count toward the elapsed time. The return value of this member function is the



elapsed time. If you call one of the elapsed member functions Chronograph::elapsed or Chronograph::elapsedHMS while the Chronograph is stopped, you will get the elapsed time when the Chronograph was stopped.

I have provided Chronograph::isstopped to indicate whether the Chronograph is stopped. It returns 1 if the Chronograph is stopped, and 0 if it is running. I recommend changing this member function to use the new bool type when it becomes available with your compiler.

Implementation Details

How can you have several timers going at once and not slow down your program? Well, you don't do it by installing hardware interrupt handling routines, although those might provide a higher resolution. You give up some resolution by simply using the timer that is already in place — the system clock.

The ANSI C clock Function

The Chronograph class uses the ANSI C clock function to get the time. The clock function returns a value of the type `clock_t`, which is the number of clock ticks since a program started. To calculate the number of seconds represented by those clock ticks, divide the `clock_t` value by the ANSI C macro `CLOCKS_PER_SEC`.

The `CLOCKS_PER_SEC` macro will likely be different from platform to platform, but the

Greg Messer works for Systems Integration and Imaging Technology, Inc. as a programmer/analyst and system administrator. He started programming 12 years ago on mainframes and now specializes in high performance client/server programming on networked PCs. Greg can be reached by e-mail at 73510.1430@compuserve.com.

equations that use it will remain the same. The CLOCKS_PER_SEC macro in Borland C is defined as 1000.0. Borland's clock function therefore returns the thousandths of a second since a program started. The MS-DOS system clock that Borland C uses in its clock

function is not really accurate to the thousandth of a second, but the Chronograph class is not meant to be that accurate either. The MS-DOS system clock is more than accurate enough for the long-duration timing tasks I had in mind.

Listing 1: Timing execution of an entire program

```
#include <iostream.h>
#include <string.h>
#include <ctype.h>
#include "chrono.h"

//allow long input lines
#define MAXLINE 1024

void getcmt(istream &in);
void usage(void);

// Global program name, for messages.
char *pgm = "GETCMT";

// -----
main(int argc, char *argv[])
{
    int i;
    double elapsed_time;

    // This Chronograph measures total
    // elapsed time of program. It starts
    // running as soon as it's instantiated.
    Chronograph chrono;
    // -----
    // Command line processing, details
    // omitted
    // ...
    // -----
    getcmt(cin);
    elapsed_time = chrono.elapsed();
    // Show the total time this program was
    // running.
    cout << endl;
    cout << pgm << " main()";
    cout << ":" Total elapsed time: ";
    cout.precision(2);
    cout << elapsed_time;
    cout << " seconds." << endl;
    cout << endl;
    return 0;
}
//End of File
```

Cross Compiler
Organon®

Choose this complete toolset: compiler, assembler, and linker all from the same company. Specially designed for embedded development and the Intel 386 EX/486/Pentium.

pSOS+ Support!

CAD-UL

**386 EX
486
Pentium**

CAD-UL, Inc.
6330-1 E. Thomas Road, #100
Scottsdale, AZ 85251
Tel. 602-945-8188, Fax 602-945-8177
eMail: sales@cadul.de
http://www.cadul.de

All trademarks owned by their respective companies.

Request Reader Service #131

- Compatible with Intel ic386.
- For use with C++ and ANSI C.
- Optional: Our powerful high-level language debugger Organon XDB 386 for ROM monitor, emulator, and RT kernels.
- Available for workstations, MS-DOS, and MS-Windows.

The operating system maintains a clock tick counter at all times, and the initial value of that counter is stored by the startup code before main is called. Subsequent calls to clock query the OS clock value, then subtract the value stored at startup. The result is the number of clock ticks since the program started.

This method of time measurement has little effect on your program's execution speed, regardless of how many Chronographs you have instantiated. The only time the Chronograph affects the speed of your program is when you call one of its member functions. The OS is servicing the clock ticks for you, whether you use them or not.

Simple Calculations

The simple calculations used by the Chronograph class depend on a few private variables. By keeping track of the clock values when the Chronograph was instantiated, when the Chronograph was stopped, and when the lap member function was last called, the Chronograph calculates elapsed

Notice To Our Subscribers

Occasionally, *C/C++ Users Journal* makes its mailing list available to vendors of products we think our readers will find interesting. Current subscribers receive free information in the mail from these vendors.

If you prefer that your name not be used in these mailings, please let us know. Just copy or clip this form and send it with your name and address to:

**C/C++
Users Journal™**
Advanced Solutions for C/C++ Programmers

**P.O. Box 52582
Boulder, CO 80322**

and lap times as the difference between two clock values.

A private member function, Chronograph::diff, accepts two `clock_t` values and returns the difference between them in seconds. This function is used whenever the elapsed or lap times need to be calculated.

Handling the Modes

As mentioned earlier, the Chronograph can be in one of two modes, running or stopped. If the Chronograph is ever stopped, it can be restarted or it can be reset and then started from zero. The elapsed and lap times can be queried regardless of the Chronograph mode. Being able to stop the Chronograph allows you to keep sections of code from being timed. For instance, you can stop the Chronograph while waiting for user input, then restart it when your program goes back to work. The time spent waiting for the user will not affect your timing statistics.

A Simple Demonstration

I show part of a filter program to demonstrate the use of the Chronograph class. The program is named GETCMT. It reads a C or C++ source file from the standard input stream and writes the comments to the standard output stream. I use this utility to give me a starting point for documenting my source code.

A simple example of the Chronograph use occurs in the `main` function of GETCMT.CPP (Listing 1). The `main` function uses this Chronograph to time the entire program. The Chronograph is instantiated a few lines down from `main`'s parameter list. It hits the ground running, so to speak. You don't need to call any member functions to start it. The program takes the elapsed time near the end of `main`, after calling the function `getcmt`. Those two lines are all that are required to time an entire program using the Chronograph.

Use `setprecision(2)` to set up your output stream before you display the elapsed time returned by the Chronograph

member functions. If you use `printf` for output, then use `%.2f` as your format string.

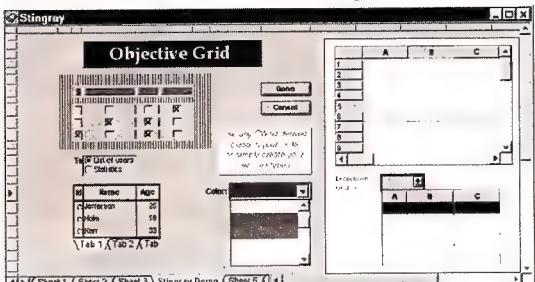
A Complex Use of the Chronograph

The `getcmt` function demonstrates a more complex use of the Chronograph class. This function uses three Chronographs — one to time the function as a whole, one to measure the time spent inside of comments, and one to measure the time spent outside of comments. Listing 2 shows portions of `getcmt`, with representative Chronograph member function calls. (Full source code is available electronically — see p. 3 for details.)

The three Chronographs are instantiated after creation of an I/O buffer and some auto variables. Since `getcmt` has yet to encounter any comments upon startup, the inside Chronograph is reset to zero shortly after instantiation, and it remains stopped. The two comment Chronographs (inside and outside) stop and start when

Tired of Wrestling with OCX/DLLs??

You need SEC++™ and Objective Grid™—two new MFC extensions that are 100% MFC compatible!



Objective Grid (pictured) is a complete grid MFC extension. You can use the grid in a CView, CWnd and even as a popup. The grid can be bound to any external data source with one virtual override. Objective Grid also features complete ODBC/DAO support, print preview, find/replace, undo/redo, UNICODE/DBCS and an extensible control architecture.

SEC++ is a group of over 40 MFC extensions including: docking views, MDI alternatives, zooming, and panning CViews, image classes for DIB/GIF/JPG/PCX/TGA/TIFF, masked edit, calendar, color well, menu button, excel-like tabbed windows, a filesystem class, encrypting and compressing CFile derivatives and much more.

Both products are VC++ 1.5x, 4.x compatible and come with full source code and no royalties. Prices: Objective Grid: \$395 SEC++: \$495. Subscriptions available.

www.stingsoft.com



"We add class to MFC!"

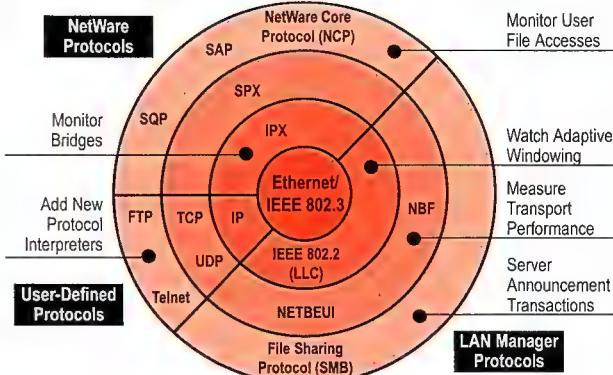
Buy now and save a bundle...

Limited time offer: SEC++/Objective Grid bundle for only **\$795!** Full source code, no royalties. Order Today! All major credit cards accepted.

□ Request Reader Service #160 □

Ethernet Protocol Analyzer The Snooper™

Become a Networking Expert with this PC-Based Analyzer for NetWare, LAN Manager and TCP/IP Networks !



The Snooper captures, decodes into multiple layers, and displays Ethernet traffic, and lets you actually add your own protocol interpreters. With source, only \$350. Ask about our EtherProbe™ !

Call us Toll-Free: 1-800-850-5755



General Software™

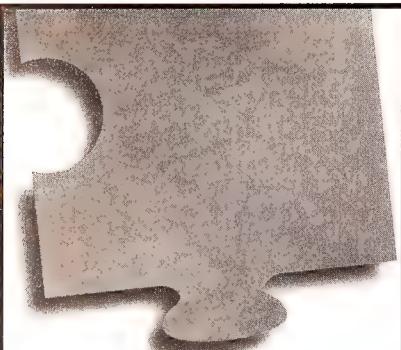
320 - 108th Ave. N.E. Suite 400 • Bellevue, WA 98004 U.S.A.

Sales: 800.850.5755 • Tel: 206.454.5755 • Fax: 206.454.5744

<http://www.gensw.com/general> • E-Mail: general@gensw.com

© 1996 General Software, Inc. All rights reserved. The GS logo, and The Snooper are trademarks of General Software, Inc.

□ Request Reader Service #132 □



**Client-Server
Relational Database System
for C/C++ Professionals**

**Fast and reliable for industrial strength applications.
Easy, no hassle, operation.**

- **Cross platform portability** on Windows and Unix
- **Client/Server** on TCP/IP
- **Multi-users** on all versions
- **Transaction support** (commit, rollback)
- **ANSI SQL**
- **Row level locking**
- Cost-based query optimizer; **no need to be an expert to write efficient SQL**

Just Logic/SQL	With client/server
SCO Unix	\$399
BSDI Unix	\$299
Windows	\$149
Windows client	\$99

Also available for Linux, Unixware, FreeBSD

**UNLIMITED RUNTIME
LICENSE INCLUDED**

30-day risk free money back guarantee. Visa, American Express, MasterCard, Eurocard accepted

Shipping: 10\$ in USA and Canada, 25\$ Express: 35\$ to other countries. All prices are in US dollars. Prices may change without notice.

(800) 267-6887
(toll-free USA and Canada)

+1(514)761-6887 (international)
(514)642-6480 (fax)

info@justlogic.com

http://www.justlogic.com

P.O. Box 63050, 40 Commerce St.,
Nun's Island QC, Canada, H3E 1V6

JUST LOGIC
TECHNOLOGIES
It makes sense!

Request Reader Service #133

Listing 2: A function that uses three chronographs

```

void getcmt(istream &in)
{
    int i, j, len;
    long linecount = 0;

    double hours, minutes, seconds;

    // Flag: 0 = not in comment,
    //        1 = in C comment
    //        2 = in C++ comment

    int incomment = 0;

    // Flag: 0 = no comment on this line
    //        1 = comment on this line

    int commentline = 0;

    char inbuff[MAXLINE + 1]; // I/O buffer

    double elapsed_time;
    double elapsed_time_inside;
    double elapsed_time_outside;

    // One Chronograph to count the time
    // spent in this function.

    Chronograph chrono;

    // Two Chronographs, to count the time
    // spent inside and outside of comments.

    Chronograph inside;
    Chronograph outside;

    // -----
    // reset the "inside comment" timer to
    // zero and stop it

    inside.reset();

    cout << endl;

    in.getline(inbuff, MAXLINE);

    while(in.eof() == 0)
    {
        linecount++;

        // demonstrate use of lap time
        if(linecount % 100 == 0)
        {
            cout << pgm;
            cout << " getcmt()";
            cout << ": 100 lines in ";
            cout.precision(2);
            cout << chrono.lap() << " seconds.";
            cout << endl;
        }

        len = strlen(inbuff);
        i = 0;

        // while chars are left on the line...
        while(i < len)
        {
            // process depending on if we are
            // in a comment and what type it is
            switch(incomment)
            {
                case 0:
                    // not inside a comment
                    // C comment detection
                    if(inbuff[i] == '/' &
                       inbuff[i+1] == '*')
                    {
                        outside.stop();
                        inside.start();
                    }
                    i += 2;
                }
            }
        }
    }
}
}

// C++ comment detection
// details omitted
// ...

case 1:
// inside a C comment

// Nested C or C++ comment
// detection details omitted
// ...

// end of C comment detected
else if(inbuff[i] == '*' &
        inbuff[i+1] == '/')
{
    inside.stop();
    outside.start();

    incomment = 0;

    i += 2;
}
else
    i++;

break;

case 2:
// inside a C++ comment
// details omitted
// ...

default:
// error, bad "incomment"
// shouldn't happen
cerr << "Error!" << endl;

return;
}

// end of switch()
}

// end of while()

if(commentline == 1)
    cout << inbuff << endl;

if(incomment == 1)
    commentline = 1;
else
    commentline = 0;

in.getline(inbuff, MAXLINE);

}

// stop all three of the timers
// and get the elapsed times
elapsed_time           = chrono.stop();
elapsed_time_inside     = inside.stop();
elapsed_time_outside   = outside.stop();

// show times spent inside and outside
// of comments. Details omitted
// ...

return;
}

//End of File

```

Listing 3: The Chronograph class

```
#ifndef CHRONO_H
#define CHRONO_H

#include <time.h>

class Chronograph
{
private:
    // when the Chronograph was started
    clock_t _start;

    // when the Chronograph was stopped
    // If _stopped != 0, the Chronograph
    // is stopped.
    clock_t _stopped;

    // when the last lap was read
    clock_t _lastlap;

    // the last elapsed time taken
    double _elapsed;

    /*
     * diff() calculates the difference IN
     * SECONDS between two clock_t values.
     */
    double diff(clock_t start, clock_t end)
    {
        return (end - start) / CLOCKS_PER_SEC;
    }

public:
    // default constructor
    Chronograph()
    {
        reset();
        start();
    }

    // destructor
    ~Chronograph()
    {}

    double start(void)
    {
        // current clock
        clock_t ct;

        // how long the Chronograph was stopped
        clock_t ctd;

        ct = clock();

        if(ct == 0)
            ct++;

        if(_stopped)
        {
            ctd = ct - _stopped;
            _start += ctd;
            _lastlap += ctd;
            _stopped = 0;
        }

        return elapsed();
    }

    else
    {
        _lastlap = ct;
        _start = ct;
        _elapsed = 0.0;

        return _elapsed;
    }

    double stop(void)
    {

```

getcmt's scanning passes into and out of comments.

When getcmt detects the start of a comment it executes the following statements:

```
outside.stop();
inside.start();
```

Likewise, when getcmt detects the end of a comment, it executes the following statements:

```
inside.stop();
outside.start();
```

That is all the code needed to maintain those two Chronographs. The simplicity of the interface is what allows you to easily add one or more Chronographs to an existing program.

getcmt stops all three Chronographs just after the end of its main while loop. Calling Chronograph::stop stops a Chronograph and also returns elapsed time.

Never Debug A Memory Leak Again.

Debugging memory leaks and premature frees is extremely time-consuming and frustrating. Even after you locate a memory leak or premature free, you still have to figure out how to fix it. And when you think you've fixed them all, your customer will use your program in a way you never tested, resulting in more leaks.

Great Circle puts an end to all that. Unlike debuggers, Great Circle doesn't simply identify memory leaks and premature frees, it eliminates them. Every single one of them. Even leaks you didn't test for. Automatically.

Automatically Eliminate Leaks Without Sacrificing Any Performance Or Control.

Simply link your C/C++ program with Great Circle and you'll eliminate the memory leaks and premature frees. All of them. Without impacting the performance of your program. Without changing the line of code. Without recompiling.

And we're not just talking new code, we're talking legacy code and third party libraries—even when you don't have the source code.

To answer some of your questions:
 1) you don't sacrifice any control or flexibility,
 2) you don't need to change your programming style,
 3) Great Circle provides a full report of the leaks it has eliminated, and
 4) Great Circle is easy-to-use. The learning curve is usually less than five minutes long.

Eliminate All Your Memory Leaks With One Phone Call.

Call 1-800-360-8388 to order Great Circle. While a maze might be an enjoyable challenge, it's no longer the best way to deal with memory leaks.

Now Also Supports THREADS and 16-BIT



Avoid Memory Leaks. Avoid The Maze.

Now, Try Great Circle Risk-Free With Our "Leak-Free" Guarantee.

"If your C or C++ program leaks or prematurely frees memory after linking with Great Circle, simply return Great Circle—within 90 days of the receipt date—and we'll refund the entire purchase price. No questions asked."

That's not simply a promise; it's our "Leak-Free" guarantee. Every word of it. No more ifs, ands, or leaks.

GreatCircle

from

Geodesic systems

800-360-8388
312-728-7196 (IL)
e-mail: info@geodesic.com

VISIT OUR WEB SITE: <http://www.geodesic.com>

□ Request Reader Service #134 □

Although not shown here, the full source code version of `getcmt` calls member function `Chronograph::elapsedHMS` to get the total time broken out as hours, minutes, and seconds. The sample program most likely will not need to measure hours and minutes, but I have some programs in daily production that use the long-range capabilities of the Chronograph. The implementation for `elapsedHMS` is available electronically with the full source code (see p. 3).

Other Uses for the Chronograph

The Chronograph has some other uses besides simple time measurement. For example, let's say you are using `Chronograph::lap` to track how long you spend on some cyclical aspect of your program, such as processing database records. If you know the total number of those records in advance you can call `Chronograph::lap` after each 1,000

MORE DEVELOPERS PROTECT.

HASP Packs More Into Less. Based on a full-custom ASIC utilizing 2500-gate, 1.5-micron E² technology, HASP packs the most advanced protection into the world's smallest key.

NSTL Study Rates HASP No. 1!*

NSTL TEST RESULTS, OCTOBER 1995[†]

Scoring Category	Aladdin HASP	Rainbow Sentinel
Security	9.3	6.3
Ease of Learning	9.1	7.1
Ease of Use	8.3	7.2
Versatility/Features	10	8.7
Compatibility	6.7	6.5
Speed of API Calls	0.9	1.2
Final Score	8.5	6.5

*For a full copy of the NSTL report, contact your local HASP distributor.

Grow With Aladdin! The fastest growing company in the industry, with over 4 million keys sold to 20 thousand developers worldwide, Aladdin is setting the standard for software security today.

1-800-223-4277
<http://www.aks.com>

ALADDIN
The Professional's Choice

North America Aladdin Knowledge Systems Inc. Tel: (800) 223 4277, 212-564 5678, Fax: 212-564 3377. E-mail: hasp sales@us.aks.com
Int'l Office Aladdin Knowledge Systems Ltd. Tel: +973-3-636 2222, Fax: +973-3-537 5796. E-mail: hasp.sales@aks.com
Germany FAST Software Security Tel: +49 89 89 42 21-37, Fax: +49 89 89 42 21-40. E-mail: info@fast-ag.de
United Kingdom Aladdin Knowledge Systems UK Ltd. Tel: +44 1753-622266, Fax: +44 1753-622262. E-mail: sales@aldn.co.uk
Japan Aladdin Japan Co., Ltd. Tel: +81 426-60 7191, Fax: +81 426-60 7194. E-mail: aladdin@po.ijinet.or.jp

Call for details of your local distributor!
© Aladdin Knowledge Systems Ltd. 1985-1996. (S)96 HASP® is a registered trademark of Aladdin Knowledge Systems Ltd. All other product names are trademarks of their respective owners. Mac & the Mac OS logo are trademarks of Apple Computer, Inc. used under license. NSTL makes no recommendation or endorsement of any product. †The NSTL report was commissioned by Aladdin.

□ Request Reader Service #135 □

Listing 3: *continued*

```

if(!_stopped)
{
    _stopped = clock();

    if(_stopped == 0)
        _stopped++;

}

return elapsed();
}

double reset(void)
{
    clock_t ct; // current clock

    ct = clock();

    if(ct == 0)
        ct++;

    _lastlap = ct;
    _start = ct;
    _stopped = ct;
    _elapsed = 0.0;

    return _elapsed;
}

/*
lap() returns the number of seconds
since lap() was last called.
*/

double lap(void)
{
    clock_t ct; // current clock
    double dl; // lap time

    if(_stopped)
        ct = _stopped;
    else
    {
        ct = clock();

        if(ct == 0)
            ct++;
    }

    dl = diff(_lastlap, ct);

    if(!_stopped)
        _lastlap = ct;

    return dl;
}

double split(void)
{
    return elapsed();
}

int isstopped(void)
{
    if(_stopped)
        return 1;
    else
        return 0;
}

/*
elapsed() returns the elapsed time in
seconds that the Chronograph has been
running.
*/

double elapsed(void)
{
    clock_t ct; // current clock

    if(_stopped)
        ct = _stopped;
    else

```

Your Roadway to Building Large Scale Systems



The Users Conference and Exhibition

Now In
its 8th
Successful
Year!

The Mecca for C++ users.

Over 2,000 serious C++ programmers will converge in Dallas on Nov. 11-15. It's the largest C++ event of the year. C++ World is created for C++ users on all platforms and levels. Meet face-to-face with the top experts and best-selling book authors. Expand and hone your C++ programming skills.

NEW and IMPROVED tracks...

Each offering vital education opportunities for the C++ programmer. Choose from:

- Methods, Analysis and Design
- Components and Frameworks
- Persistence
- Distribute Objects
- Software Development
- Large Scale Design

PLUS: a special "Gurus Only" accelerated track specifically designed for advanced programmers with 3+ years C++ experience.



Past and Present Exhibitors (partial list):

Addison Wesley • Alsys • Borland • C++ Report • C/C++ Users Journal • Cayenne Technologies • CenterLine • Datalink • DevelopMentor • Epersoft • IBM • ICON Computing • IEEE Computer Society • InterSolv • JOOP • Java Report • McCabe & Associates • Metrowerks • MultiQuest • Neuron Data • Object Design • Object International • Object Magazine • ObjectSpace • O'Keefe & Assoc • Pencom • POET Software • Prentice Hall • ProtoSoft • Pure Software • Rogue Wave • Semaphore • SIGS Publications • Software Emancipation • Symantec • TakeFive Software • Teknekron Comm. Systems • TRW • UniSQL • VERITAS Software • The X Journal • XVT Software • Zinc Software

What You Will Learn:

Build on your C++ foundation through a variety of educational classes — over 40 to choose from. You'll learn:

- How to select a distributed object technology path by comparing and contrasting ActiveX, CORBA and OpenDoc.
- How to move from C++ to Java™.
- The secrets to reuse.
- How to get the most out of Visual C++ and MFC.
- How to design and manage class libraries.
- Multithreading tips in C++.
- How to decouple and leverage code.
- Effective exception handling.
- Testing & debugging tips.
- Understand and use the standard template library.
- How to use modular design with C++.
- Managing pointers and complexity in C++.
- Using patterns effectively.
- Using ODBMS with C++.
- The Unification of methods.
- Web programming tips and techniques.

November 11-15, 1996

Grand Kempinski Hotel
Dallas, Texas

For more information
Call 212.242.7515

Reach our website at
<http://www.sigs.com>

Email us at conferences@sigs.com

Sponsored by:
C++REPORT

Presented by:
SIGS
CONFERENCES

records and use those lap times to dynamically update an estimated completion time.

I also put the Chronograph to good use in a "spinner" class that I wrote to indicate a program's progress. By repeatedly cycling through a series of characters ('/', '\', '|'), I display a character that appears to be spinning.

The spinner class contains a private Chronograph that makes sure the screen update does not happen too often. No matter how frequently an application tries to

update the spinner, it only displays itself at one second intervals.

Conclusion

I use the Chronograph class daily. Often I am asked how long it will take to execute some program. With as little as two lines of code I can add a Chronograph to any C++ program and get the timing statistics I need. With only a few more lines of code I can create any timing statistic I can imagine.

The need for timing exists in virtually every significant program. With this Chronograph, you can add timing capabilities in a few moments and then focus on the more interesting parts of the program. That's the real power of a great tool. □

Listing 3: continued

```
{
    ct = clock();
    if(ct == 0)
        ct++;
}

_elapsed = diff(_start, ct);
return _elapsed;
}

double elapsedHMS(double &Hours,
                  double &Mins,
                  double &Secs);
};

#endif

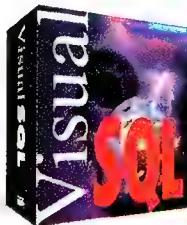
// end of chrono.h

//End of File
```



**"With
Visual SQL,
I can
quickly
develop
Client/Server
applications in
C++ without leaving
Microsoft Visual C++."**

Build industry standard Client/Server applications with Visual SQL® inside Visual C++™. Visual SQL turns Microsoft® Visual C++ into a full-blown Client/Server development environment. Visual SQL integrates seamlessly with Microsoft Developer Studio using intuitive wizards to generate MFC Client/Server code that is fully ODBC-compliant. Visual SQL automates and shortens every stage of the development cycle and offers developers the rapid application development of Visual Basic® and the power of compiled C++ code.



CALL NOW TOLL FREE FOR MORE INFORMATION

1-888-831-6222

<http://www.blue-sky.com>

(outside US 619-459-6365)

□ Request Reader Service #136 □

**The
CUG
Directories
Volumes
II, III, & IV**

**\$10 each
All three for \$28**

Each volume contains detailed descriptions of public domain and shareware C source code. Every file is indexed by appropriate keywords, filenames, and authors. Also included are extensive reviews extracted from the documentation and *C/C++ Users Journal* articles.

Volume II covers disks 200-249;
Volume III covers disks 250-299;
Volume IV covers disks 300-349

Call, FAX or write

**The C Users' Group
1601 West 23rd St., Suite 200
Lawrence, KS 66046
913-841-1631
FAX 913-841-2624**

The Java and C Connection

Sure, Java is a neat new toy. But it's nice to know you can mix in a bit of C code, from time to time, to beef up a Java program.

In recent months, one of the hottest technologies to hit the streets has been Java. It has taken the information systems industry by such a storm that most software vendors have jumped on the bandwagon to develop products for Java. Meanwhile, many companies still have large investments in applications written in other languages, such as C. My goal here is to show you how to interface to C from Java. I assume that you have a working knowledge of both Java and C.

Java provides a native method interface which can be used to invoke functions written in other languages. Currently only an interface to C is supported. (Of course, C++ source code can be accessed via C functions). There are approximately six steps required for connecting a Java program to C. The two main steps include writing the Java program and C functions. The remaining steps are required for gluing the two together.

Here is a simple example to illustrate each of these six steps. The example consists of a Java program that calls a C function to print the message "Hello Java" (a twist on the typical "Hello World" message). To keep the first example simple, no parameters are passed to the C function and no values are returned from it. Figure 1 provides a visual representation of the various components created from the following six steps required to make the Java and C connection:

- 1) Write a Java program.
- 2) Compile it.
- 3) Generate a header file for the C function.

Anil Hemrajani currently provides software engineering and training consulting services to a Fortune 500 corporation in McLean, VA. He can be contacted at anil@patriot.net or via <http://www.patriot.net/users/anil/>.

- 4) Generate a stub file for the C function.
- 5) Write the C function.
- 6) Build a dynamic link library with the C function.

Let's look at each of these steps in more detail.

Step 1

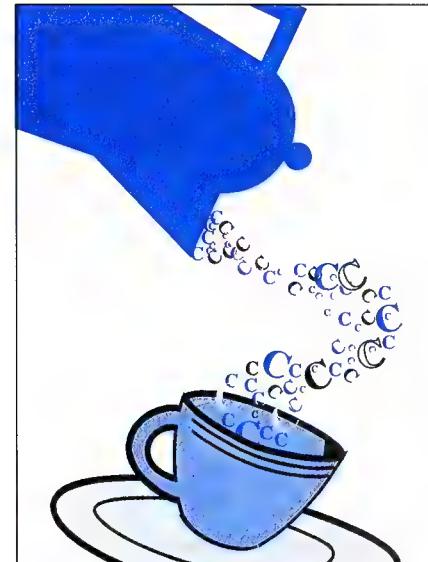
The following is a simple Java program that defines a native method `SayHi`. Its implementation is provided in step 5:

```
public class HelloJava
{
    public native void SayHi();

    public static void
        main(String args[])
    {
        System.loadLibrary("hello");
        new HelloJava().SayHi();
    }
}
```

Native methods are declared much the same way as regular Java methods, with two differences. The declaration must contain the `native` keyword, and there is no body for the method in Java. The body is provided in the C function. The `loadLibrary` method (part of the Java `System` class) must be called prior to using the native method so that the dynamic link library containing the C function can be loaded. The `new` expression must be used to instantiate a native method.

What I have just shown here is one way of instantiating a native method. Alternatively, the native method could be declared in another Java class. In that case, we instantiate that Java class using a `new` expression, then simply call the native method off of the class. This style is used in Listing 3 and described below.



Step 2

Compile the Java class using a Java compiler, such as `javac` utility provided with the Java Development Kit (JDK). Running the following command will generate a Java byte code file, `HelloJava.class`:

```
javac HelloJava.java
```

Step 3

Generate a header file for the C function using the `javah` utility. The generated header file (`HelloJava.h`) basically provides a prototype for the C function, which is covered in step 5.

```
javah HelloJava
```

Step 4

Generate a stub function, using the `javah` utility, to connect the Java class to the C function. The following command will generate a stub file named `HelloJava.c`:

```
javah -stubs HelloJava
```

Step 5

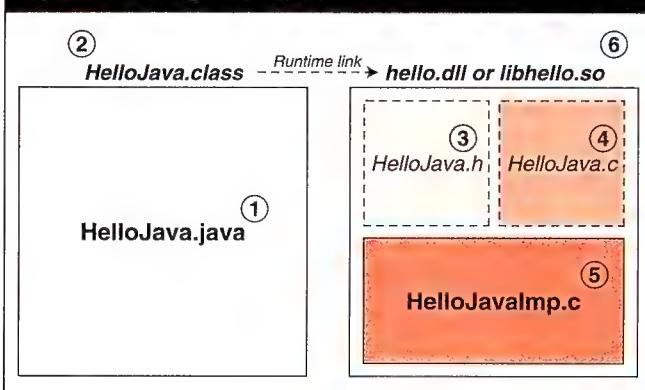
Provide an implementation for the C function, `SayHi`:

```
/* HelloJavaImp.c: Implementation of
   Java native method, SayHi(). */
#include "HelloJava.h"
#include <stdio.h>

void HelloJava_SayHi(struct HHelloJava *javaObj)
{
    printf("Hello Java!\n");
}
```

Three things in the above source code deserve mention. First, the header file generated in step 3, HelloJava.h, must be included in the C module. Second, notice the name of the C function,

Figure 1: Using C with Java



VICTOR
Image Processing Library

Speed and power processing for all popular image formats

Victor Image Processing Library
DOS \$199, 16-bit DLL \$299,
32-bit DLL \$499

- ★ Fast load & save BMP, TIFF, PCX, GIF, TGA, JPEG
- ★ Powerful image processing: brightness, contrast, sharpen, smooth, equalize, create filters, resize, rotate, +++
- ★ Accurate color reduction to optimum, specific, or standard palette
- ★ Print halftone, diffusion scatter, or color
- ★ Scan with all HP scanners

Visit our home page for demos, pictures, and sample code
<http://www.catenary.com/victor>

Call or fax to order... voice: 314-962-7833 fax: 314-962-8037

Catenary Systems
470 Bellevue St Louis MO 63119

No royalties Source code available visa/mc/cod.

□ Request Reader Service #137 □

HelloJava_SayHi. It contains the name of the Java class and the native method separated by an underscore. You will come across this naming convention (*class_method*) whenever you are interfacing with native methods in Java.

Last, an automatic parameter is always passed as the first parameter to a native method. It is essentially a handle to the Java class invoking the native method. You can think of this parameter as the *this* keyword in C++. I discuss how to use this parameter below.

Step 6

Build a dynamic link library. Here is a simple makefile which contains targets for building a dynamic link library on Windows 95, using the Microsoft C/C++ compiler, and a shared library on Solaris. The environment variable JAVAHOME points to the root directory of the JDK:

```
all:
    @echo Please specify a target: Win95 or Sun.

Win95:
    cl HelloJavaImp.c HelloJava.c \
    -I$(JAVAHOME)\include -I$(JAVAHOME)\include\win32 \
    -Fehello.dll -MD -LD $(JAVAHOME)\lib\javai.lib

Sun:
    cc -G -I$(JAVAHOME)/include -I$(JAVAHOME)/include/solaris \
    HelloJavaImp.c HelloJava.c -o libhello.so
```

Once you have completed the above six steps successfully, you are ready to run the sample Java and C application. Here is how you would test the application and the output you should receive:

```
C:\ANIL\Doc\Articles\Java-C\Example1> java HelloJava
Hello Java!
```

Passing Parameters and Returning Values

Now that I have shown you the basic steps required to connect Java programs to C functions, it is time to get a little fancier by passing parameters to a C function and returning values from it. Since steps 2, 3, 4 and 6 are similar for all cases, I concentrate on steps 1 and 5 which deal with a Java program and C function(s), respectively. But before plunging into the specifics of these examples, I provide a quick overview of Java data types and how they correspond to C data types.

The following table provides a list of Java native data types and their machine sizes:

Type	Size
byte	8-bit
char	16-bit
short	16-bit
int	32-bit
float	32-bit
long	64-bit
double	64-bit

As you might notice, some Java data types are larger than their equivalent C data types. For example, a char is 16 bits in Java and

(typically) eight bits in C. The numeric data types are similar to those you find on UNIX systems, but they are typically larger than ones on MS-DOS based systems. An int is usually two bytes on MS-DOS, for example, but it is four bytes in Java. The sizes for Java data types are the same across all supported platforms, which makes the language more portable.

Listings 1 and 2 show a sample Java program and a sample C function (tested on Windows 95 and Solaris), respectively. The purpose of this sample application is simple. The Java class calls a C function with certain parameters. The C function prints the values in these parameters to stdout using printf and returns the number of characters printed to the Java class, which in turn prints a message indicating how many characters were printed.

Let's review some of the significant points in these examples. In the Java program, notice how I placed the native method in a separate class (PrintInC) from the "main" class (ThePrinter). I did this to show you an alternative way of invoking native methods. Also notice the static block in the PrintInC class. This is a good place to automatically perform tasks when the class is loaded, which in my case happened to be loading a dynamic library:

```
static { System.loadLibrary("print"); }
```

Now, let's consider the C implementation. First notice that I include a StubPreamble.h header file. This is a JDK include file which contains necessary structures (such as HArrayofInt), functions prototypes (such as makeCString), and macros (such as unhand). Note also the mapping of data types used in Java versus what I had to use in C:

```
public native int doPrint(long l,
    int i[],
    int count,
    double d,
    String s);

long PrintInC_doPrint(struct HPrintInC *this,
    int64_t l,
    HArrayofInt *ai, long iCount,
    double d,
    struct Hjava_lang_String *s)
```

Notice how each int has been mapped to long. An int may be only two bytes on MS-DOS but is always four bytes in Java. Similarly, a Java long argument has been mapped to the type int64_t in C. The other two notable parameters in these examples are a Java array (int[]) and a Java object (String). The Java int[] array gets mapped to the C structure HArrayofInt.

HArrayofInt is one example of a standard naming convention used in the StubPreamble.h header file. Other examples include HArrayofLong, HArrayofByte, and so on. Java objects also use standard naming conventions such as H for handle, the name of the language (e.g. java), the package (class library) containing the class (lang), and the name of the object (String). So a structure for the Java String object, which is part of the lang package, is named Hjava_lang_String.

Finally, let me point out a C macro (unhand) and function (makeCString) used in this example. The unhand macro provides

access to the variables in Java's C Structures, such as the HArrayofInt, as shown here:

```
i = unhand(ai)->body;
```

Another example of unhand is shown in the next section.

The makeCString function takes a Java String as an argument and returns a corresponding char pointer. There are two other functions related to conversions between Java strings and C character arrays, makeJavaString and java2CString. The former is covered in the next section along with an example. The latter is similar to makeCString, except that it works more like strncpy — it copies a Java String object into an existing char array as shown in the following example:

```
char OutFile[127+1];
javaString2CString(pOutFile, OutFile, sizeof(OutFile));
```

Accessing Java Objects

So far I have shown you how to call C functions from Java. Now I show you how to go the other way, accessing Java objects from C. Java provides the ability for C functions to access data members in Java objects, create instances of Java classes, and invoke dynamic and static methods in the Java class. Listings 3 and 4 demonstrate how to accomplish most of these tasks.

Listing 3 contains mostly the same type of Java source code as my previous examples, but one thing is done slightly differently and

C and C++ DOCUMENTATION

C-METRIC™ (\$59) – Complexity/Quality

- Calculates "cyclomatic" path complexity for functions and system
- Counts lines with comments, code, and 'C' statements

C-CALL™ (\$69) – Function Hierarchy

- Tree-Diagram showing function hierarchy
- Table-of-Contents of functions versus files
- Summary and detailed cross-reference of functions

C-CMT™ (\$69) – Function Comment

- Generates and inserts function comment blocks
- Can be re-run to update the comment blocks

C-LIST™ (\$69) – Lists or Reformat

- Action-Diagrams show logic/control flow
- Reformats source to various standardized formats

C-REF™ (\$69) – Cross-References Identifiers

- Local/global/define/parameter summary or cross-reference
- Produces class-hierarchy tree-diagram for C++ classes

C-BROWSE™ (\$free in C-DOC) – Windows Tree Viewer

- Graphically view C-CALL function-trees or C-REF class-trees

C-DOC™ (\$199) – DOS/Windows Package (\$395 value)

- All 5 programs integrated as 1 overall C-DOC program
- Processes multiple directories/files up to 10,000 lines
- Unconditional 30-day money-back guarantee of satisfaction

C-DOC™ Professional (\$299) – DOS, Windows, OS/2

- All features of C-DOC, processes 1,000,000 lines, 3-ring binder/case

!! NEW 6.0 !! Windows Graphic-Tree Viewer

SOFTWARE BLACKSMITHS INC.

6064 St Ives Way
Mississauga ONT
Canada L5N-4M1

Voice/Fax **(905) 858-4466**
<http://swbs.idirect.com>

CALL NOW!

Listing 1: Sample Java program

```
// ThePrinter:
// Call native methods to print stuff

public class ThePrinter
{
    public static void main(String args[])
    {
        int count=4;
        int i[] = new int[count];
        i[0] = 10;
        i[1] = 75;
        i[2] = 95;
        i[3] = 115;

        int printed = new PrintInC().doPrint(
            25,
            i, count,
            100.33,
            "Hello C");

        System.out.println("Java: " +
            printed +
            " chars printed");
    }
}

class PrintInC
{
    public native int doPrint(long l,
        int i[],
        int count,
        double d,
        String s);

    static
    { System.loadLibrary("print"); }
}
// End of file
```

Converting FORTRAN to C?

We've got the right tools to ease you through this seemingly daunting process!

FOR_C, our robust source code conversion software will take legacy FORTRAN (with support for VAX and other extensions), and automatically generate readable, maintainable ANSI C code.

FOR_STRUCT, our FORTRAN restructuring tool, will clean up your unstructured code, transforming it into clean, manageable code. Whether you are translating to C, or maintaining in FORTRAN, your code will be highly improved!

Call or fax today for more information on these and other FORTRAN tools. Or check out our web site for downloadable demos and product data!

COBALT BLUE

555 SUN VALLEY DRIVE
SUITE K-4
ROSWELL, GA 30076, USA
TEL: (770) 518-1116
FAX: (770) 640-1182
INTERNET: <http://www.cobalt-blue.com>



□ Request Reader Service #139 □

worth mentioning. The JavaAccessor class illustrates how a Java class containing multiple native methods can be instantiated and then individual methods in that class can be invoked, as shown here:

```
JavaAccessor ja = new JavaAccessor();
..
System.out.println("PATH=" + ja.getEnv("PATH"));
ja.deleteFile("dummy.txt");
```

Listing 4 requires a little more explanation, as there are several new things I have not covered yet. First, the getEnv native method returns a Java object (String), unlike the other examples I have shown so far which return native data types. Notice how you have to use makeJavaString, the opposite of makeCString, to return a Java String object:

```
return makeJavaString(value, strlen(value));
```

You can access data members from the “automatic” parameter that is passed to every native method, in this case the pointer to HJavaAccess. You simply reference a data member in a Java class by using the unhand macro to dereference the Java object:

```
long JavaAccess_deleteFile(struct HJavaAccess *javaObj,
                           struct Hjava_lang_String *fileName)
{
    ...
    unhand(javaObj)->delRC=rc;
```

CQL++

C++ B-tree, SQL, ODBC file handler

Includes complete C++ source.
Portable to any environment with
an ANSI C++ compiler.

All the advantages of SQL and
ODBC without sacrificing the
efficiency of B-tree level access.

Single user and client/server
configurations.

Royalty free for embedded
applications!

ANSI 1989 Level 2 SQL with
extensions.

ODBC driver allows use of GUI
interface tools.

Transaction processing with
commit, rollback, and automatic
recovery from hardware and
software failures.

Scorable cursors, fetch previous,
persistent cursors.

Single user.....\$495.00
Client/Server.....\$995.00

Machine **I**ndependent **S**oftware **C**orporation
1819 Ridgemont Columbia, MO. 65203 USA
Voice: +1.314.446.4242 Fax: 446.4243
E-Mail: machine@sba.cerfnet.com CompuServe 76256,2743
<http://sba.cerf.net/~machine>

□ Request Reader Service #140 □



C++ AND WINDOWS PROGRAMMING

COMING SOON
TO:

WASHINGTON, DC

May 13 - 17

NEW YORK

May 20 - 24

BOSTON

June 3 - 7

TORONTO

June 10 - 14

SEATTLE

June 17 - 21

SALT LAKE CITY

July 29 - August 2

COLUMBUS

September 9 - 13

DALLAS

September 23-27

MINNEAPOLIS

October 7 - 11

CHICAGO

October 14 - 18

DENVER

November 4 - 8

ATLANTA

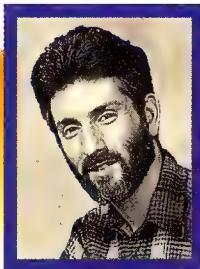
November 18 - 22

**CALL US TODAY TO
FIND OUT MORE!**

PHONE: (415) 905-2702

FAX: (415) 905-2222

E-MAIL: sdconfs@mfi.com

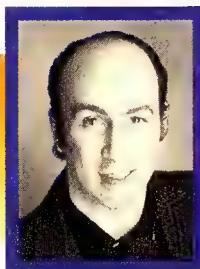


Taught by Dan Saks

C++ Fundamentals Not compiler-specific;
suitable for any programming platform

- Learn today's premiere programming language from a leading expert and teacher
- New! Integrated Labs to master concepts through practical programming exercises (bring your own computer or laptop)
- Progress step-by-step through the key language features

- Understand how data abstraction, inheritance and polymorphism support sound design practice
- Move from C to C++ in just 2 days
- Prepare for Windows-based programming with C++



Taught by Richard Hale Shaw

**Windows 95/Windows NT Programming
with Visual C++ 4.0 and MFC 4.0**

- Learn the fundamentals of Windows programming with VC++/MFC 4.0
- New! Integrated Labs (bring your own computer or laptop)
- Explore the MFC architecture
- Use VC++/MFC 4.0 for Rapid Application Development
- Apply Visual techniques to build, derive, and modify new C++ classes
- Learn MFC's document-view architecture
- Build dialog-only applications

- Form views, toolbars with cool user-interface components
- Use VC++ 4.0's built-in Component Gallery objects
- Learn how to back-port to Windows 3.1
- Use OLE controls in your VC++ applications
- Create MFC extension DLLs
- Use MFC to get a jump-start on OLE
- Discover the wonder of OLE Automation
- Build Automation servers

FOR MORE CLASS INFORMATION:

[HTTP://WWW.MFI.COM/SDCONFS](http://WWW.MFI.COM/SDCONFS)

Produced by the Software Development Conference & Show Group of Miller Freeman, Inc.
600 Harrison Street, San Francisco, CA 94107

Listing 2: PrintInCImp.c: C functions for Java class

```
#include <stdio.h>
#include <StubPreamble.h>
#include "PrintInC.h"

/** Java to C demo: Print data values ***/
long PrintInC_doPrint(struct HPrintInC *this,
    int64_t l,
    HArrayOfInt *ai, long iCount,
    double d,
    struct Hjava_lang_String *s)
{
    int charsPrinted=0, idx=0;
    long *i;

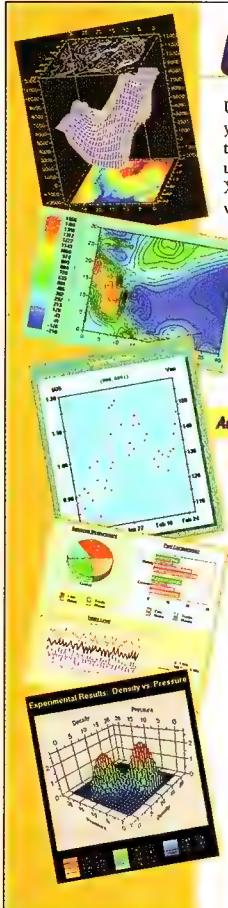
    /* Print value of "long" (Java int) */
    charsPrinted += printf("C: l = %ld\n", l);

    /* Print array of longs */
    i = unhand(ai)->body;
    for (idx=0; idx < iCount; idx++)
        charsPrinted += printf("C: i[%d] = %d\n", idx, i[idx]);

    /* Print value of "double" */
    charsPrinted += printf("C: d = %f\n", d);

    /* Print Java string */
    charsPrinted += printf("C: s = %s\n", makeCString(s));

    /* Return total characters printed */
    return charsPrinted;
}
// End of file
```



Ultimate Charting

Unprecedented performance and control – whether you develop for Windows or Motif, KL Group has the ultimate charting control for you. Loaded with unique features, Olectra Chart for Windows and XRT widgets for Motif are powerful enough to build virtually any 2D or 3D chart or graph:

- X-Y Plots, Bar Charts, Pie Charts, Combination Graphs, Financial Graphs and Logarithmic Scientific Charts
- Real-time Updates to Build Scrolling Strip Charts
- Built-in Zooming, Scaling, Rotation
- Chart Labels and Annotations Anywhere
- Tunable Auto Precision Axis
- Dynamic Time & Date Axis on Any Plot

Advanced technology that out-performs the competition

Windows: With distinct 2D/3D OCXs and DLLs, Olectra Chart is more than 4X faster than First Impression 2.0 and Graphics Server SDK 4.0 in standard mode redraw speed. Olectra Chart includes toll-free support and a money-back guarantee – at only \$249!

Motif: XRT/graph and XRT/3d are the world's leading Motif charting widgets. These award winning components can be used on over 18 UNIX platforms and can be easily integrated with all popular GUI builder tools. Call today to get your free 30-day evaluation!



The Leader in GUI Components

1-800-666-4723

Tel: 416.594.1026 Fax: 416.594.1919
info@klg.com www.klg.com

* Price in US dollars. Prices may be higher outside North America.
XRT is a registered trademark and Olectra Chart is a trademark of KL Group Inc.
All other product names are trademarks or registered trademarks of their respective owners.

□ Request Reader Service #141 □

Finally, you can instantiate Java classes and invoke methods within them. Three C functions accomplish this — execute_java_constructor, execute_java_dynamic_method, and execute_java_static_method.

execute_java_constructor (shown below) expects at least four parameters. It can require more, depending on the number of parameters a constructor requires. The first parameter indicates the language/execution environment, which for our purposes is always zero for Java. The second parameter is the name of the Java class. The third is currently reserved for future use. The fourth is the signature of the constructor, and the remaining arguments are any parameters to pass to the constructor.

The signature specified in the fourth parameter is made up of a set of parentheses which contain symbols for a method's parameters. For example, since JavaAccessor's constructor expects no parameters, it simply has the signature "()":

```
hJavaAccessor=(HJavaAccessor *)
    execute_java_constructor(0, "JavaAccessor", 0,
        "()");
```

Similarly, execute_java_dynamic_method (shown below) also expects at least four parameters. It can require more, depending on the number of parameters the Java method requires. The first parameter indicates the language environment, which again is zero for Java. The second parameter is a reference to the Java object created via the execute_java_constructor function. The third parameter is the name of the method to invoke. The fourth is the method's signature, and the remaining arguments are any parameters to pass to the method.

As shown in this example, the Java method "printRC" expects a single argument of type int (I) and does not return any values, so it

Listing 3: Demonstrates Java Object Access

```
public class JavaAccessor
{
    public synchronized native String getEnv(String var)
        throws java.lang.NullPointerException;

    public native int deleteFile(String fileName);
    public int delRC=0;

    public static void main(String args[])
    {
        System.loadLibrary("ja");
        JavaAccessor ja = new JavaAccessor();

        try System.out.println("PATH=" +
            ja.getEnv("PATH"));
        catch(java.lang.NullPointerException e)
        {
            System.out.println("getEnv returned NULL!");
        }

        ja.deleteFile("dummy.txt");
        System.out.println("Return code from" +
            " deleteFile() = " +
            ja.delRC);
    }

    public void printRC(int RC)
    {
        System.out.println("Java: RC=" + RC);
    }
}
// End of file
```

has a V (for void) to indicate this at the right side of the closing parenthesis:

```
if (hJavaAccessor)
    execute_java_dynamic_method(0,
        (HObject *)hJavaAccessor,
        "printRC",
        "(I)V", rc);
```

One good way to find out what letters to use for data types in the signature is to define native methods in a dummy java class and then generate a stub file using the javah utility. You can then inspect the generated .c file, which will have comments indicating the signature types. For example, listing 5 shows a stub file generated from the following Java class:

```
class test
{
    public native long LONG();
    public native float FLOAT();
    public native boolean BOOLEAN();
    public native void MIX_PARMS(
        int i, long l,
        String s, byte b,
        float f, double d,
        char c, boolean bool,
        short si, int[] ai);
}
```

Exceptions and Thread Synchronization

C functions called by Java classes can throw exceptions, which can be caught by Java methods. This is accomplished by the `SignalError` C function provided in the Java library. Consider this line taken from Listing 4:

```
SignalError(0,
    "java/lang/NullPointerException",
    0);
```

Java is a multi-threaded language, hence C functions used in Java could potentially be accessed concurrently by various threads in a given Java program. Java provides a data-locking mechanism which permits thread-safe operations in Java and native methods. For C functions written for Java, thread-safe operations can be accomplished by using the `synchronized` keyword in the Java method declaration (as shown below) and three C functions provided in the JDK: `monitorWait`, `monitorNotify`, and `monitorNotifyAll`.

```
public synchronized native String getEnv(String var);
```

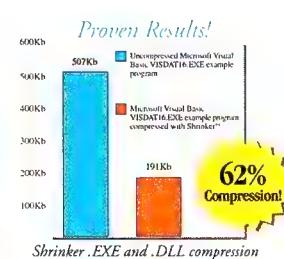
Conclusion

Java is a feature-rich language that comes bundled with packages for everything from local file I/O, to socket communications, to GUI programming, to classes for the World Wide Web. Still, there might be times where you need to drop down to C or C++ to accomplish specific tasks, or perhaps to leverage existing C routines or a C API library. For these times, it is nice to know that Java provides the ability to do so. However, one thing to keep

"Invisible Shrinking" for your Windows® .EXEs and .DLLs!

Cut your distribution and storage costs in half, improve network performance and protect your valuable software from decompilers with this innovative new technology. Any 16-bit Windows .EXE or .DLL compressed with **Shrinker™** will transparently decompress itself into memory at runtime, so nobody will ever know it has been "shrunk"!

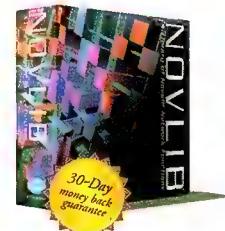
Compatible with Visual Basic™, Delphi™ C++, CA-VO® and *much more*, **Shrinker** compresses virtually all your 16-bit Windows programs by 50% or more and there are no runtime royalties to pay. If you need more disk space or distribute programs electronically, you need **Shrinker** today!



Make all your programs NetWare®-aware the easy way!

The only NetWare library for Visual Basic™, Delphi™, C/C++, CA-VO and CA-Clipper, this award-winning library features over 450 network functions to build comprehensive NetWare support into your programs.

If you write Windows, DOS or DOS extended programs for Novell networks, you need **NOVLIB**!



Save time, memory and disk space!

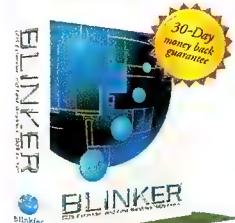
Share code and data between Windows and DOS, easily update or replace parts of your program and simplify the migration of code to Windows using **Blinker's** new DLL feature.

Blinker® feature checklist:

- unique ability to compress Windows and DOS extended EXEs and .DLLs by 50% or more
- fastest available 16 & 32 bit Windows linker
- DOS extender to directly access up to 16Mb of memory, fully compatible with DPMI, VCI and XMS programming standards and featuring a comprehensive API
- Windows and DOS hosted linker also creates .DLLs for DOS extended programs
- unique dual mode feature to run the same .EXE automatically in either real or protected mode
- a memory efficient super "SPAWN" command to run large programs, such as a word processor or a report generator, from inside your main program with less than a 10Kb memory overhead
- Microsoft C/C++, FORTRAN, Pascal, Assembler, Borland C/C++, Assembler, Symantec C++, Watcom C/C++ and CA-Clipper support
- full support for CodeView, Soft-ICE and Periscope debuggers
- no runtime royalties

The fastest available 16- and 32-bit Windows linker and the only product to transparently compress Windows and DOS extended programs by 50% or more, **Blinker** is also a royalty-free DOS extender, allowing DOS programs to directly access up to 16MB of memory.

Call today to see why over 75,000 programmers already use this award-winning product.



Special Offer!

Shrinker™ \$149.99!
NOVLIB® \$299.249!
Blinker® \$299.249!

To order or for your FREE demo disk
Call: 1-804-747-6700
Fax: 1-804-747-4200
URL: <http://www.blinkinc.com>

the Standard for Development!

Blinkinc
8001 W. Broad St., Richmond VA 23294

In Europe contact: Blinkinc Ltd.

TEL: +44 1222 712444 • FAX: +44 1222 700888

Professional Real-Time Tools

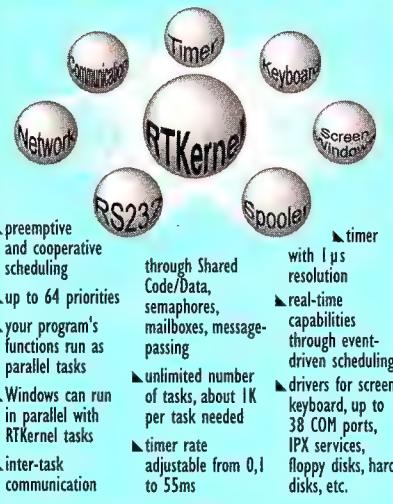
RTKernel Real-Time Multitasking for DOS and 16-Bit Embedded Systems

Real-Time Multitasking for:

• Borland C/C++, Microsoft C/C++, or Borland Pascal

RTKernel supports:

MS-DOS 3.0 or higher, Embedded Systems without DOS, Paradigm Tools (info available), Turbo Debugger, CodeView



- ▶ preemptive and cooperative scheduling
- ▶ up to 64 priorities
- ▶ your program's functions run as parallel tasks
- ▶ Windows can run in parallel with RTKernel tasks
- ▶ inter-task communication
- ▶ Developer's License: \$550
complete Source Code: add \$500
no run-time royalties

RTTarget-32 Cross Development System for 32-Bit Embedded Systems

32-Bit Development System for:

• Borland C/C++, Microsoft C/C++, Watcom C/C++

RTTarget-32 supports:

Intel 386/486 EX/SX/DX(2), as little as 16K RAM/ROM

RT TARGET 32

- ▶ boot code for floppy, hard disk, EEPROM disk
- ▶ optional paging and RAM remapping
- ▶ FLAT memory model with physical = logical addresses
- ▶ Developer's License: \$1700
complete Source Code: add \$1000
no run-time royalties
- ▶ interrupt latency <5µs
- ▶ program download at 115200 baud
- ▶ supports standard PCs and controller boards
- ▶ fully supports the C/C++ run-time systems (printf, malloc, etc.)
- ▶ locator for Win32 PE-files

On Time MARKETING
Professional Programming Tools

Free demo disk! <http://www.on-time.com>
Please request Info Kit C
<ftp.on-time.com>
info@on-time.com

In North America,

please contact:

On Time Marketing
88 Christian Avenue
Setauket, NY 11733

USA
Phone (516) 689-6654
Fax (516) 689-1172

Other countries:

On Time Marketing

Karolinienstrasse 32

20357 Hamburg

GERMANY

Phone +49-40-437472

Fax +49-40-435196

email 100140.633@compuserve.com

□ Request Reader Service #143 □

in mind is that, whenever possible, any C functions written for Java should be kept as portable and ANSI compliant as possible. Portability is a big benefit of Java, and it would be nice if you could have the C code be portable as well.

If you do not yet own a book on Java, try downloading the free documentation

from Sun Microsystems' FTP site (<ftp.javasoft.com>). *The Java Tutorial*, by Mary Campione and Kathy Walrath, is an excellent source of information for learning Java. It is available on the Web site in two formats, HTML and Postscript. □

Listing 4: Access to Java environment from C

```
/** Delete a local file ***/
long JavaAccessor_deleteFile(
    struct HJavaAccessor *javaObj,
    struct Hjava_lang_String *fileName)
{
    HJavaAccessor *hJavaAccessor;
    int rc=unlink(makeCString(fileName));
    unhand(javaObj)->delRC=rc;

    hJavaAccessor=(HJavaAccessor *)
        execute_java_constructor(
            0,
            "JavaAccessor",
            0,
            "(())");

    if (hJavaAccessor)
        execute_java_dynamic_method(
            0,
            (HObject *)hJavaAccessor,
            "printRC",
            "(I)V",
            rc);

    else
        printf(
            "Unable to create hJavaAccessor\n");
}

return rc;
}
```

Listing 5: Stub File generated from Java Class

```
#include <StubPreamble.h>

/* Stubs for class test */
/* SYMBOL: "test/LONG()", Java_test_LONG_stub */
stack_item *Java_test_LONG_stub(stack_item *_P_,struct execenv *_EE_) {
    Java8 _tval;
    extern int64_t test_LONG(void *);
    SET_INT64(_tval,_P_,test_LONG(_P_[0].p));
    return _P_ + 2;
}
/* SYMBOL: "test/FLOAT()", Java_test_FLOAT_stub */
stack_item *Java_test_FLOAT_stub(stack_item *_P_,struct execenv *_EE_) {
    extern float test_FLOAT(void *);
    _P_[0].f = test_FLOAT(_P_[0].p);
    return _P_ + 1;
}
/* SYMBOL: "test/BOOLEAN()Z", Java_test_BOOLEAN_stub */
stack_item *Java_test_BOOLEAN_stub(stack_item *_P_,struct execenv *_EE_) {
    extern long test_BOOLEAN(void *);
    _P_[0].i = (test_BOOLEAN(_P_[0].p) ? TRUE : FALSE);
    return _P_ + 1;
}
/* SYMBOL: "test/MIX_PARMS(IJLjava/lang/String;BFDCZS[I]V", Java_test_MIX_0005fp
ARMS_stub */
stack_item *Java_test_MIX_0005fpPARMS_stub(stack_item *_P_,struct execenv *_EE_) {
    Java8 _t2;
    Java8 _t7;
    extern void test_MIX_PARMS(void *,long,int64_t,void *,long,float,double,
long,long,long,void *);
    (void) test_MIX_PARMS(_P_[0].p,(_P_[1].i),GET_INT64(_t2,_P_+2),(_P_[4].p),(_P_[5].i),(_P_[6].f),GET_DOUBLE(_t7,_P_+7),(_P_[9].i),(_P_[10].i),(_P_[11].i),(_P_[12].p));
    return _P_;
}
/* End of file */
```

Floating-point Summation

Addition may be one of the first operations you master when learning mathematics, but it is probably one of the last ones you learn how to do safely on a computer.

Introduction

The summation of large quantities of floating-point numbers is a simple task that arises frequently in practical applications. Some of the most common applications are numerical integration and calculating averages. But this task turns out not to be as straightforward and predictable as it sounds.

First let's look at the obvious way. Writing a routine to sum a series of numbers may seem trivial. Listing 1 is my version of a Simple Summation routine. It uses a local `float` variable `sum` to accumulate a running total of all elements in the array `flt_arr`. This routine is simple, straightforward, and frequently gives bad results.

The problem is that addition of floating-point numbers is different from addition of integers. Integer addition is like doing addition by hand — you line up two numbers one over the other and add one digit at a time, starting from the least significant. Of course computers do this in binary. Floating-point addition includes a step like this too, but first the numbers must be shifted so that corresponding digits line up. But if the two numbers have sufficiently different exponents then some least significant digits will have to be discarded.

For example, if we have a base ten system with one digit of accuracy then valid values include 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 20, 30, 40, 50, 60, 70, 80, 90, 100, 200, and so on. In this system, $10.0 + 1.0 = 10.0$ because 11.0 is not a valid value. In such a system, 10.0 would indeed be the best possible answer for the sum of these two numbers, so there is no problem. But if three or

more numbers are to be summed then errors accumulate, and the order of the additions becomes important.

If, for another example, there are not enough bits of mantissa to resolve a difference of one part in ten billion, then:

$$\begin{aligned} (1.0 + 1.0e10) + -1.0e10 \\ = 1.0e10 + -1.0e10 \\ = 0.0 \end{aligned}$$

but

$$\begin{aligned} 1.0 + (1.0e10 + -1.0e10) \\ = 1.0 + 0.0 \\ = 1.0 \end{aligned}$$

In this example, all the information from one summand was lost because two of the summands were much larger than the other. Usually when you are summing large sets of floating-point numbers there will not be such pronounced differences between individual terms, but errors will grow from the accumulation of many small errors, and from the difference between the running total and the individual terms.

An obvious fix is to add numbers with higher precision. If we start with an array of `float` we should then keep our internal result as a `double`. If we started with an array of `double` we should then keep our result as a `long double`.

Listing 2 is almost like Listing 1 but with the type of sum changed to `double`. This method gives much higher precision than Simple Summation, and is an excellent approach, when available. But if calculations are already being performed in the highest precision available then extending precision is not an option.

Order in the Court

Much of the imprecision stems from the order of operations. Precision is lost most quickly when the magnitude of the two num-



bers to be added is most different. With summation, one of the operands is always the sum so far, and in general it is much larger than the numbers being added to it. This effect can be minimized by adding the numbers in order from the smallest to the largest. (Smallest in magnitude, that is. A very negative number is still large in the sense that its least significant bit has little precision. Bits will be discarded in adding -0.001 to 1000.0 but not in adding -1000.0 and 1000.0.)

This effect can be minimized by sorting the numbers by magnitude before summing. Listing 3 does this using `qsort` and a comparison function. First it makes a local copy of `flt_arr` in `lcl_arr`, then it sorts this array with `qsort`. Finally, it sums the sorted terms with a running total like the other methods shown so far.

The results are still not good. Sorting is a much slower operation than addition, and the sum still quickly becomes much larger than the numbers being added, so accuracy is not greatly enhanced.

In fact, particular details of the numbers being summed critically impacts how well sorting works. The best case for sorting is where the sum is near zero and the distribution is nearly symmetric. These conditions will keep the sum near the same magnitude as the numbers being added. These conditions are satisfied in test cases 2 and 3 in the test program (Listing 8), but rarely in real life.

Sorting does bring predictability. Sorted Summation always gives the same answer for the same set of numbers, even if those numbers are presented in different orders. This is not entirely true if you use the

Evan Manning has degrees in Applied Physics from Caltech and Stanford. He has been working as a self-taught C programmer in defense and space applications for the past 8 years. Currently he works for Telos Information Systems as a consultant at NASA's Jet Propulsion Laboratory. He can be reached at manning@alumni.caltech.edu.

comparison function in Listing 4. A number and its negation compare equal, and so may show up in any order in the sorted sequence. When you need predictability, use a comparison function like that in Listing 5. This comparison function breaks such ties to attain complete predictability.

Listing 1: Simple Summation

```
float f_add(float * flt_arr)
{
    long i;
    float sum = 0.0;

    for (i = 0; i < ARR_SIZE; i++)
        sum += flt_arr[i];
    return sum;
} /* End of File */
```

Listing 2: Extended Precision method

```
float fx_add(float * flt_arr)
{
    long i;
    double sum = 0.0;

    for (i = 0; i < ARR_SIZE; i++)
        sum += flt_arr[i];
    return sum;
} /* End of File */
```

Another approach to keeping the operands of each addition similar in magnitude is to add in pairs. Listing 6 presents a routine that loops through the array adding pairs of numbers and reducing the size of the array by half each iteration. At each addition, the summands represent the same number of original entries, and so should have similar magnitudes.

Pairwise Summation proves to be both faster and more accurate than Sorted Summation, and further refinements are possible which would improve it further. But it still cannot compete with ...

The Right Way

An even better way to sum many numbers is Kahan Summation [1]. With this method, a correction is maintained along with the current running total, effectively increasing the precision. Listing 7 implements Kahan Summation. Each term is first corrected for the error that has accumulated so far. Then the new sum is calculated by adding this corrected term to the running total. Finally, a new correction term is calculated as the difference

between the change in the sums and the corrected term.

Algebraically, this routine seems to do nothing special. Substituting the line:

```
new_sum = sum + corrected_next_term;
```

into:

```
correction
= (new_sum - sum) -
corrected_next_term;
```

gives:

```
correction
= ((sum +
corrected_next_term) - sum) -
corrected_next_term;
```

which would simplify to:

```
correction
= corrected_next_term -
corrected_next_term
= 0;
```

for exact numbers. But for real floating-point numbers, the correction term will frequently be nonzero. Using it does improve accuracy. This method is more complicated than the other methods shown here, but that complexity is easily hidden in a function.

Here is an example program. It is organized around a pair of arrays,

```
float flt_array[ARR_SIZE];
double dbl_array[ARR_SIZE];
```

The main routine (Listing 8) runs four tests. For each test the code repeatedly fills these arrays (OUTER_ITER repetitions) and calls the routine loop to shuffle and sum the values.

- The first test fills the arrays with uniform random values ranging from 0.0 to 1.0. I use four calls to rand for each element.

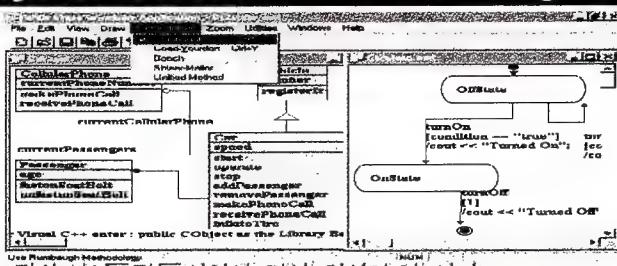
Listing 3: Sorted Summation

```
float fs_add(float * flt_arr)
{
    long i;
    float lcl_arr[ARR_SIZE], sum = 0.0;

    for (i = 0; i < ARR_SIZE; i++)
        lcl_arr[i] = flt_arr[i];
    qsort(lcl_arr, ARR_SIZE, sizeof(float),
          flt_abs_compare);
    for (i = 0; i < ARR_SIZE; i++)
        sum += lcl_arr[i];

    return sum;
} /* End of File */
```

PRESENT YOUR **O-O** DESIGNS **WITH CLASS!** Object-Oriented C++ Design Tool



*** NEW ADD-ON: REVERSE JAVA ***

- ◆ Reverse Header files with EASE!
- ◆ Generate .h and .cpp files
- ◆ **Reverse Engineer** C++, JAVA, Delphi & ODBC Databases
- ◆ Scripting interface generates **Custom Reports** and **Custom Code**
- ◆ **OLE Support** In-place editing
- ◆ Supports Borland C++, 3.X, 4.X, Visual C++1.X, 2.X, 4.X
- ◆ 350 Page Tutorial - "**O-O Modeling, C++ Programming**"
- ◆ **Template driven** Code and Report generation for **State, Class** and **Object Interaction** diagrams
- ◆ Customize **color** and **fonts**
- ◆ **Generate Code** in C++, Java, Delphi, Pascal, VFoxpro, VB, SQL
- ◆ Supports Unified & all leading **UML** methods

16 and 32 Bit **WITH CLASS** pricing starts at \$295

Call for your **FREE** Evaluation Copy Today! www.micrgold.com

MICRGOLD SOFTWARE, INC. 76 Linda Lane Edison, NJ 08820 v.908.668.4779 f.908.668.4386

□ Request Reader Service #144 □

This assures me of at least 60 bits of randomness even when rand produces only 15 bits (the minimum permitted by the ANSI/ISO standard).

- The second test is like the first except that the arrays are rescaled so that they run from -1.0 to 1.0.
- The third test uses random values with a bell-shaped distribution.
- The final test is a numerical integration to estimate the value of pi by integrating one quarter of a circle.

loop applies all five summation methods on each of the two data sizes (float and double), and then shuffles the arrays and repeats for SHUFFLE_ITERS iterations.

Results

Over a large number of runs a pattern becomes obvious: Simple and Sorted Summation are tied for least accurate, Pairwise Summation is a little better, Kahan Summation is better still, and the champion is Extended Precision, as shown in Listing 2. (I also provide a function, fpx_add, to do Extended Precision Kahan Summation,

which is not shown, but is available on the code disk and via ftp. This function is just like the regular precision fk_add, except that the local floating-point variables are defined as doubles, instead of floats. See page 3 for downloading details.)

In my test on one machine I found Simple Summation to be consistently the fastest. Extended Precision was a bit faster than Simple Summation when Simple Summation used float and Extended Precision used double, but was about 50 times slower when Simple Summation used double and Extended Precision used long double. Apparently this platform favors double over float and does not support

Listing 4: Simple comparison

```
int flt_abs_compare(const void * a,
                     const void * b)
{
    float fa = fabs(*((float *)a));
    float fb = fabs(*((float *)b));
    if (fa == fb)
        return 0;
    else if (fa < fb)
        return -1;
    else
        return 1;
} /* End of File */
```

long double in hardware but simulates it in software.

Sorted Summation was typically about 100 times slower than Simple Summation, rendering it unsuitable for most applications. Because sorting scales as $n \log(n)$ (where n is the number of operands) while all other methods discussed here scale as n , sorting becomes less suitable for larger data sets. Pairwise Summation is about 3-4 times slower than Simple Summation with doubles, and so would be a good bet if it

Listing 5: Tie-breaking comparison

```
int flt_abs_compare(const void * a,
                     const void * b)
{
    float va = *(float *)a;
    float vb = *(float *)b;
    float fa = fabs(va);
    float fb = fabs(vb);
    if (fa < fb)
        return -1;
    else if (fa > fb)
        return 1;
    else if (va < vb)
        return -1;
    else if (va > vb)
        return 1;
    else
        return 0;
} /* End of File */
```

New! CSIM18!

A SIMULATION TOOLKIT:

- New features
 - confidence intervals
 - run-length control
 - execution monitor
- More platforms
 - WIN95, Power Mac
- Use C/C++ to build process-oriented models of complex systems
- Multiple platforms
- Low cost solution

CALL NOW FOR FREE DEMO DISK



Mesquite Software, Inc.

3925 West Braker Lane • Austin, TX 78759-5321
U.S. 1-800-538-9153 or 1-512-305-0080

FAX 1-512-305-0009

e mail: info@mesquite.com

WWW: <http://www.mesquite.com/>

□ Request Reader Service #145 □

Don't sweat the small stuff, just ship it!

How to ruin your reputation for quality

Wasted time and re-work from interruptions caused by any exception can make the most loyal customer think of ways your software might have handled it better. This is the worst possible time for your exception process to create further problems, imply disrespect for the customer, reveal unprofessionalism, and/or suggest that the headline above may summarize your policy on quality.

Unlike error handlers, which are named functions or macros, *exception processes* aren't named units of code. This makes them nearly impossible to locate and test; and makes it easy to rationalize ignoring them as "small stuff".

The XMan Exception Management System is the only tool available that is designed especially to enable software professionals to find and remove these time bombs from multi-file projects coded in C. A well-written book (over 350 pages) describes how to use the system, and also offers techniques, examples, a tutorial, and some commentary on writing programmed messages. XMan runs under DOS 3.1 and later; no wasteful GUI is needed. A single-user license fee is \$1195. Sanders-Indev, Inc. 5 Seneca Ave., Geneseo, NY 14454-9508. 716-243-2091.

Free demo and more information at

<http://www.sanders-indev.com>

□ Request Reader Service #146 □

were not eclipsed by Kahan Summation. Kahan Summation turns out to be about half as fast as Simple Summation when using double precision.

Simple Summation, Extended Precision, and Kahan Summation all use very little extra storage space. As implemented here, Sorted Summation uses an additional array the same size as

Listing 6: Pairwise Summation

```
float fp_add(float * flt_arr)
{
    long i, j, limit;
    float sum[ARR_SIZE / 2 + 1];

    if (ARR_SIZE == 2)
        return flt_arr[0] + flt_arr[1];
    else if (ARR_SIZE == 1)
        return flt_arr[0];

    for (i = j = 0; i < ARR_SIZE / 2; i++, j += 2)
        sum[i] = flt_arr[j] + flt_arr[j + 1];
    if (ARR_SIZE & 1)
        sum[ARR_SIZE / 2] = flt_arr[ARR_SIZE - 1];
    limit = (ARR_SIZE + 1) / 2;

    while (limit > 2) {
        for (i = j = 0; i < limit / 2; i++, j += 2)
            sum[i] = sum[j] + sum[j + 1];

        if (limit & 1)
            sum[limit / 2] = sum[limit - 1];
        limit = (limit + 1) / 2;
    }

    return sum[0] + sum[1];
}
/* End of File */
```

VSMT-OS
Virtual State Machine Technology

Feature rich
Real-time Multi-tasking Operating
Systems

Up To 50% saving on development costs
Get your Real-Time, On-Time!

500 task switches/second - 2ms scheduler slot
Typical timings achieved from common
processor configurations

Probably the best, and fastest OS for small
embedded systems available today!

Complete designs and implementations for
most 8,16 and 32-bit processors from just \$20!

See our Web page!

Available soon, total Windows front-end
configuration!

PO Box 2542, Wimborne, Dorset, BH21 3YE ENGLAND
Tel.+44 (0)1202 691501 Email.101543.401@compuserve.com
Web. <http://ourworld.compuserve.com/homepages/bbt>

□ Request Reader Service #147 □

the original and Pairwise Summation uses an array half the original size. But in some applications the summation would be the last operation performed on the array and so the same storage could be reused for the sorted array or the pairwise partial sums. Alternatively, these algorithms could be optimized somewhat to reduce storage if needed.

Mix and Match

Various combinations of the algorithms discussed here may also prove useful. On platforms like mine, where float operations are always slower than double operations, you should always sum with double precision even when the summands are float. This gives you higher precision and greater speed. Kahan Summation with all local variables of type double would be faster and more accurate than the version of Kahan Summation shown here.

Similarly, on platforms where long double operations are not more expensive than double operations, any of these methods should be performed with long double.

Even though sorting is slow and does not improve accuracy, it does guarantee to give the same sum for different shufflings of a given data set. When you need this predictability, sort first and then perform Kahan Summation or some other method.

Kahan Summation is always a good bet if you don't have specific timing information available for your target platform. If you can do timing on the target platform, then explore the possibility of combining Extended Precision and Kahan Summation.

Final Thoughts

It seemed natural to write the code for this article in C instead of C++ because C++ features might prove a distraction for some readers and because the code is relatively simple. To my surprise, I found the lack of some of my favorite features of C++ jarring. I particularly missed // comments, function templates, reference parameters, and the ability to mix declarations in with my code.

I recommend Reference [1] to all readers; it really does deliver "What Every Computer Scientist Should Know About Floating-Point Arithmetic" (But Was Afraid to Ask). □

References

- [1] David Goldberg. "What Every Computer Scientist Should Know About Floating-Point Arithmetic," *ACM Computing Surveys*, Vol. 23, #1, March 1991, pp. 5-48.

Listing 7: Kahan Summation

```
float fk_add(float * flt_arr)
{
    long i;
    float sum, correction, corrected_next_term, new_sum;

    sum = flt_arr[0];
    correction = 0.0;
    for (i = 1; i < ARR_SIZE; i++)
    {
        corrected_next_term = flt_arr[i] - correction;
        new_sum = sum + corrected_next_term;
        correction = (new_sum - sum) - corrected_next_term;
        sum = new_sum;
    }
    return sum;
}
/* End of File */
```

Listing 8: Tests performance of summation methods

```

int main(void)
{
    float flt_arr[ARR_SIZE];
    double dbl_arr[ARR_SIZE];
    double val;
    double x;
    int j, k;
    long i;

    /* seed random number generator from time */
    srand((unsigned int)time(0));

    /* initialize arrays for test 1: uniform random 0.0 - 1.0 */
    printf("\n% Simple Extended"
          " Sorted Pairwise Kahan\n", "Unif 0-1 ");
    for (j = 0; j < OUTER_ITER; j++)
    {
        for (i = 0; i < ARR_SIZE; i++)
            flt_arr[i] = dbl_arr[i] =
                (rand()*(double)RAND_MAX*RAND_MAX*(double)RAND_MAX
                 + rand()*(double)RAND_MAX*RAND_MAX
                 + rand()*(double)RAND_MAX
                 + rand())
        / ( RAND_MAX*(double)RAND_MAX*RAND_MAX*(double)RAND_MAX
             + RAND_MAX*(double)RAND_MAX*RAND_MAX
             + RAND_MAX*(double)RAND_MAX
             + RAND_MAX);
        loop(flt_arr, dbl_arr, "Unif 0-1 ");
    }
    print_and_reset_minmaxerr();

    /* reinitialize for test 2: uniform random -1.0 - 1.0 */
    printf("\n% Simple Extended"
          " Sorted Pairwise Kahan\n", "Unif -1-1 ");
    for (j = 0; j < OUTER_ITER; j++)
    {
        for (i = 0; i < ARR_SIZE; i++)
            flt_arr[i] = dbl_arr[i] =
                (rand()*(double)RAND_MAX*RAND_MAX*(double)RAND_MAX
                 + rand()*(double)RAND_MAX*RAND_MAX
                 + rand()*(double)RAND_MAX
                 + rand()) * 2.0
        / ( RAND_MAX*(double)RAND_MAX*RAND_MAX*(double)RAND_MAX
             + RAND_MAX*(double)RAND_MAX*RAND_MAX
             + RAND_MAX*(double)RAND_MAX
             + RAND_MAX)
        - 1.0;
        loop(flt_arr, dbl_arr, "Unif -1-1 ");
    }
    print_and_reset_minmaxerr();

    /* reinitialize for test 3: more or less Gaussian */
    /*
     * Simple but not very accurate way to make a gaussian: just
     * sum lots (here 12) uniform random variables.
     */
    printf("\n% Simple Extended"
          " Sorted Pairwise Kahan\n", "Gaussian ");
    for (j = 0; j < OUTER_ITER; j++)
    {
        for (i = 0; i < ARR_SIZE; i++)
        {
            val = 0.0;
            for (k = 0; k < 12; k++)
                val += rand();
            flt_arr[i] = dbl_arr[i] = val / (double)RAND_MAX - 6.0;
        }
        loop(flt_arr, dbl_arr, "Gaussian ");
    }
    print_and_reset_minmaxerr();
}

```

```

/* reinitialize for test 4:
 * numerical integration to calculate pi */
/*
 * The area of a circle is pi * radius^2. When the radius is
 * 1.0 the area is pi. So we approximate pi as 4 times the
 * area of one quarter of the unit circle, and get the area
 * of that quadrant using numerical integration.
 */
printf("\n% Simple Extended"
      " Sorted Pairwise Kahan\n", "Pi ");
for (j = 0; j < OUTER_ITER; j++)
{
    for (i = 0; i < ARR_SIZE; i++)
    {
        x = (1 + 2 * i) / (double)(2 * ARR_SIZE);
        flt_arr[i] = dbl_arr[i] =
            sqrt(1.0 - x * x) * (4.0 / ARR_SIZE);
    }
    loop(flt_arr, dbl_arr, "Pi ");
}
print_and_reset_minmaxerr();

/*
 * Omitted from listing: code to print times to nearest tenth
 * of a second
 * ...
 */

return EXIT_SUCCESS;
}
/* End of File */

```

DON'T WORRY ABOUT EXPENSIVE SOFTWARE DOCUMENTATION...

The screenshot shows the MyFriend programming workbench. At the top, there's a menu bar with File, Edit, Insert, Views, Goto, Project, Window, Options, Help. Below the menu is a toolbar with icons for New, Open, Save, Print, etc. The main window has a title bar 'MyFriend' and a status bar at the bottom. A large tree view on the left shows a class hierarchy under 'C/C++'. A search dialog box is open in the center, titled 'Search function name'. It contains fields for 'Function name' (with 'C/C++' highlighted), 'File' (set to 'C:\MyFriend\MyFriend\MyFriendView.h'), 'Level' (set to '1'), and a list of member functions. At the bottom of the dialog are OK, Cancel, Help, and Function name buttons.

C/C++

new feature: Project Management System

... use *MyFriend*, the programming workbench for source code and NASSI-SHNEIDERMAN CHARTS.

Programming and documentation in one tool

MyFriend Release 2.0 \$575 + shipping
Visa, MasterCard, Amex and Diners accepted

GSK

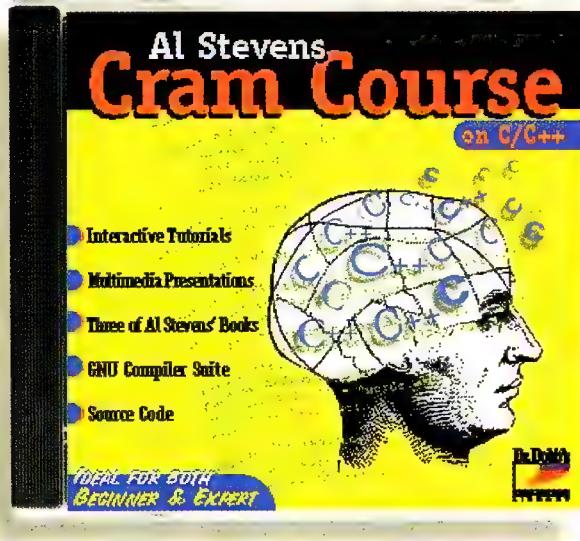
GSK GmbH Phone +49 911 63 75 85
Unt. Pfaffensteigstr. 71 FAX +49 911 63 75 86
D-91126 Schwabach BBS +49 911 63 24 968
Germany CompuServe 100414,3044

□ Request Reader Service #148 □

Learn C/C++ Programming

Presenting the Al Stevens Cram Course on C/C++

THE C/C++
PROGRAMMERS
ESSENTIAL
RESOURCE

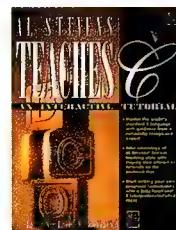


Includes these Books by Al Stevens:



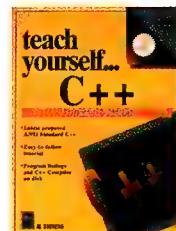
Welcome to Programming

Straight forward explanations of computer files, code, and language makes programming simple.



Al Stevens Teaches C

The perfect introduction to C with an organized, structured approach, teaching good programming through good example.



Teach Yourself ...C++

This book demonstrates the hows and whys of C++ program design in an accessible, informative style.



SYSTEM REQUIREMENTS:

- 486/40 (486/66 recommended)
- SVGA
- CD-ROM
- Mouse
- 8 MB RAM
- Win95 or NT
- Sound Card (Not essential)
- 5 MB Disk Space
- (Will run on Windows 3.1, but requires manual installation)

Complete this form and return to :

Fax: 913-841-2624; E-mail: orders@mfi.com;
Int'l: Use mail, fax, e-mail, or call 913-841-1631

Please send me an *Al Stevens Cram Course on C/C++* CD-ROM today! My price is \$69.95 (plus \$2.00 U.S./Canada; all other countries \$10.00). Please add sales tax in the following states: CA (8.5%); GA (6%); IL (6.25%), KS (6.9%), NY (8.25%), TX (8.25%)

Name _____

Address _____

City _____

State _____

ZIP _____

Country _____

Phone _____

Check

E-mail address _____

VISA

Mastercard

Amex

Credit Card# _____

Exp. Date _____

Signature _____

C0996

Order Today! 800-548-1971

Full Money-Back Guarantee! No Questions Asked!

In Search of a Portable Screen Library

If you thought finding a good screen library was difficult, you're right. But author Michaels supplies some useful hints to aid your search.

This article is for those of you starting a new project and finding yourself in the unavoidable position of trying to locate a nice, simple, cross-platform user interface library. I've spent many hours over the past few months in just such a search, checking books, magazines and the Internet. I'm hoping this will save you a great deal of time if you find yourself in a similar position.

Two Extremes

The most basic screen I/O available in C++ is the iostreams library. Personally I'd be happy to just use cout and cin in all my programs, if only I had not run into some drawbacks. One problem is that not all compilers support cin and cout correctly, especially if they create programs for Microsoft Windows. Later versions of these compilers may correct this problem, but I must work with what I have. Another major concern is that many users want a more standard user interface, one that looks like the rest of the programs that run under their operating system. Trying to create a user interface that includes menus, edit boxes, push buttons, and the like using just cout- and cin-style commands would take an enormous amount of work. That's the point of trying to locate the screen library in the first place — to avoid that work.

Going to the other extreme, there are plenty of commercial cross-platform libraries available. But, for a variety of reasons — among them cost and complexity — I quickly ruled them out. Lots of software magazines run some good comparison articles if you think this is the route you want to take.

Having ruled out those two possibilities, I began to explore the options between the two extremes. The biggest difficulty I face in my

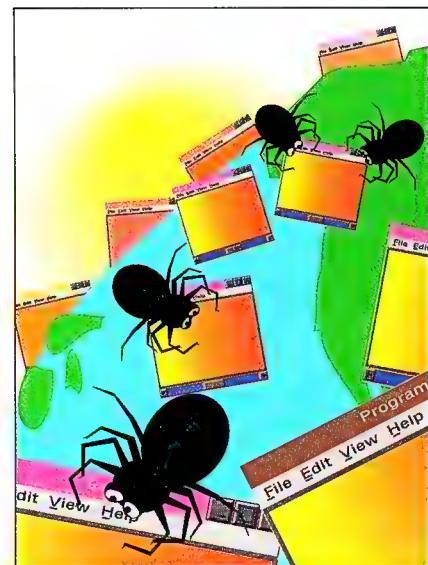
applications is that I want my user interface library to work on character-based systems (like MS-DOS) as well as graphical systems (like Windows). It's easy to find systems for one or the other category, but finding one that supports both is almost impossible.

Character-Based Libraries

Curses

Let's start with character-based user interfaces. One of the most well-known character-based screen libraries is curses (a.k.a. ncurses). If you have programs written in curses or are willing to work with the curses API, you can find several ports of curses available for different systems. The ncurses library works on any POSIX-conformant UNIX platform, and can be obtained by ftp from most GNU collections. (For example, see Item #1 in the box "Starting Places.") A note about downloadable software: whenever retrieving libraries, especially from the Internet, be sure to check that you have the latest stable version. Also, since information on the Internet can change daily, if you can't access a reference, try locating the library or program with a search engine. If you need a DOS port of curses, there's a nice version available from the SIMTEL archive called PDCURSES. (For more information, see Item #2.) It also works on OS/2 and an X11 version is available. There are also ports of curses available for Windows NT and other platforms.

If you find, as I did, that curses doesn't give you enough high-level functionality, you can check for libraries that work on top of curses. A couple of good ones are versions of WXWIN or TK. However, I found out the hard way that just because screen libraries work with curses and you have a DOS version (or a version for whatever system you're on), it does not mean the screen library will work on your system. A lot of libraries that work with curses have other UNIX dependencies in their code as well.



More functionality

I did finally manage in my search to turn up a good number of character-based libraries with higher functionality than curses. They are available for a nominal fee. Several are located at the Simtel archives under the msdos/cpluspl directory. There is a library called Sword written for djgpp compiler users. It works with other compilers too. You can check the djgcc collections, such as those under the SIMTEL archive for more information. (See Item #3.)

I also found a very comprehensive library, C/Windows Toolchest, available from Tom Mix software (Item #4). All of these character-based options are reasonably priced. With most of them, you can try before you buy. Another option is Turbo Vision, which Borland freely includes with its C++ compilers. These make a good alternative if you feel very comfortable with the API and you don't need support for other platforms. However, I do need support for other platforms and compilers, so I decided to continue my search.

D-Flat

One last character-based choice I found was the library D-Flat. It is written in C and chronicled in a series of articles in Al Stevens' "C Programming" column of *Dr. Dobb's Journal* [1]. D-Flat is freely available. There are several versions floating around, including a C++ version, an OS-2 version and a djgpp version. However, the latest and most stable version I could turn up was version 20 (Item #5). You can also get a C++ wrapper class (Item #6) to use with version 20. D-Flat 20 works on a DOS

Laura Michaels is a senior software engineer at INTERCOMP, creator of entertainment, hobby, and custom software. She can be reached through INTERCOMP, section INT-IP, Box 6514, Delray, FL 33482 or <http://members.aol.com/lauram3017/index.html>.

Starting Places

Here is a list of screen libraries mentioned in this article, and places where you can either find them, or more information about them. The Internet changes daily, and we can only supply ftp and web site addresses that were current as of this writing. Most ftp sites have readme files to direct you to new sites if the locations change. Also, you can visit the author's web page at <http://members.aol.com/lauram3017/index.html> for updated information.

1. ncurses	ftp://prep.ai.mit.edu/pub/gnu/ncurses-1.9.9e.tar.gz
2. PDCURSES	http://www.gu.edu.au/gwis/the/markh.html or ftp://grind.isca.uiowa.edu/pc/simtel/msdos/screen/ \ pdcurs22.zip
3. Sword	http://mimine.iut.univ-metz.fr/~borysgr/sword.web/home.html ftp://wuarchive.wustl.edu/systems/msdos/simtel/ \ vendors/djgpp/vtk/sw200a.txt
4. C/Windows Toolchest	Tom Mix Software 1132 Commerce Dr., Richardson, TX 75081 +1-214-783-6001. http://www.mixsoftware.com
5. D-Flat	ftp://ftp.mv.com/pub/ddj/packages/dflt20.zip
6. D-Flat wrappers	ftp://ftp.mv.com/pub/ddj/packages/dfwrap.zip
7. wxWindows	http://www.aiai.ed.ac.uk/~jac/s/wxwin.html
8. YACL	http://www.cs.sc.edu/~sridhar/yacl.html
9. TCL/TK	http://www.sco.com/Technology/tcl/Tcl.html or http://www.sunlabs.com:80/research/tcl/
10. Java	http://www.javasoft.com/
11. PIGUI FAQ	http://www.zeta.org.au/~rosko/pigui.htm

platform and can be compiled as a 16-bit or 32-bit library for Borland, Microsoft, or Watcom. (Although I could not find specific mention of Symantec support in the latest version, it shouldn't be too hard to get a Symantec version running with support for so many other compilers built in.)

Graphical Libraries

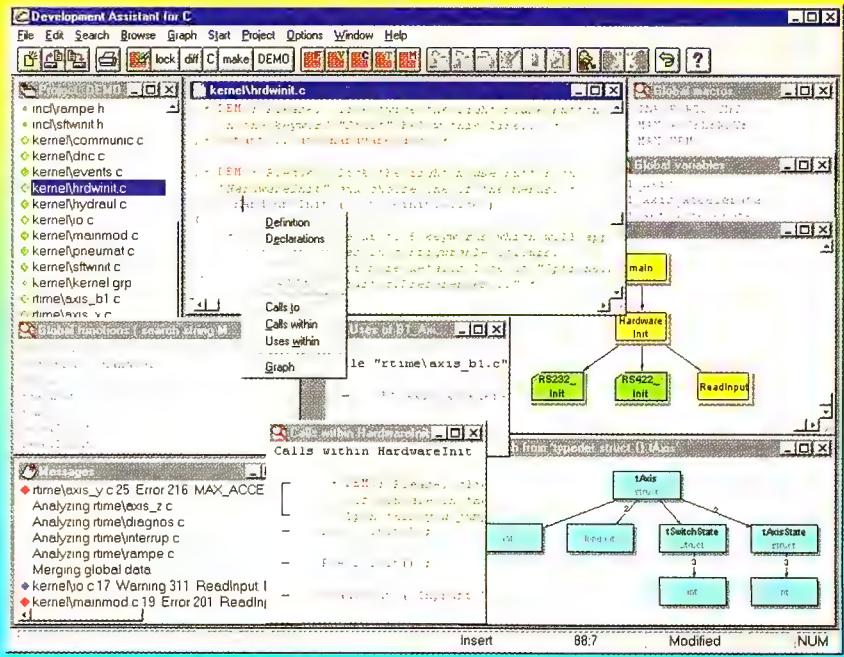
Moving on to the graphical end of the spectrum, I located several comprehensive libraries freely available over the Internet. The first one I found was wxWindows. (See Item #7.) wxWindows supports or has ports in progress for Windows, UNIX, Macintosh, OS/2, Nextstep, VMS, and Amiga. I also found some references to YACL. (See Item #8.) YACL supports Windows and OS/2 platforms and UNIX systems with Motif. One nice thing about YACL: there is a YACL book [2], and it's available in local bookstores. Be aware that YACL has some strings attached if you use it for commercial purposes.

Continuing with freely available libraries, I'd be remiss if I failed to mention TCL/TK. (See Item #9.) TCL often runs as an interpreted language, but the libraries for it are all

Development Assistant for C or PL/M

available in 16- and 32-Bit Vers. for Windows

Directly supported compilers: ANSI, Keil, IAR, Organon, Borland and Microsoft, PL/M 51 and PL/M *86



- The Static Analyzer (with complete parser for several dialects of C) that goes beyond compilers' per-module limits and detects usage conflicts, dead code, syntax errors... Generate metrics, make-dependencies and database for extremely accurate browsing...

- The Editor with a source highlighting, fully configurable keyboard (Brief and Borland em.), hiding, macro programming. Just point and click on any symbol in source and jump to one of its browse or graph views... Start any other application from command line entered in a comment (e.g. WinWord with a description, problem report or requirements definition file).

- Browsers for the definition, declarations, usage of selected symbol; global symbols (incremental search), value-change and function calls, usage or calls within function. Jump to adequate editor line or graph for any symbol. Search and replace throughout the whole project.

- Versatile call-hierarchy-graphs and type-graphs with zooming, editing, grouping, colouring (public, static, lib, group), mosaic-printing, hyper-jumps to Editor, Left-right or top-down.

- Powerful and easy to understand Macro language for the integration of DOS or Windows programmes (Compiler, Make, Debug, VCS) and viewing the results in a Browser or Editor.

With DA you can think and act on the project-level; everything is available and surveyable on (right) mouse-click or key-press. DA is especially worthwhile on large projects, software reviews or for the introduction of new staff members.



Web: <http://www.ristancase.ch/da>

RistanCASE GmbH
Zielackerstr. 19
CH-8804 Wallisellen
E: info@RistanCASE.ch
Tel. +41 1 833 07 57
Fax +41 1 833 06 14

□ Request Reader Service #149 □

written in C. In fact, TCL was originally designed as a library to be used in C programs. Some libraries also exist to ease the embedding of TCL functions in C or C++ code. TK is the actual screen interface library, but it requires the TCL library to run.

TCL and TK are available on a variety of platforms including Windows (3.1, 95, and NT), UNIX, Macintosh, and OS/2. TK includes both standard GUI versions and a curses version to support character-based applications. The UNIX version is very well supported. Rewrites for Windows and Macintosh are currently in the works and may not be as stable. You can also find a number of recent books about this library in bookstores [3] [4].

Have Some Java

Some would say that Java, with its Abstract Window Toolkit (AWT) is the ultimate cross-platform language. I found some references to the AWT libraries. (See Item #10) The AWT is Java's official screen library. It's similar to C++'s iostreams in that it allows screen I/O across all platforms, but has more functionality in creating GUIs. I personally would love to see something like this for C++, as it would automatically solve my portability problems. The Java syntax is very similar to C/C++, so if you want to be absolutely sure that you're using a portable screen interface, switching to Java to write your code may be a very attractive alternative.

Of course, as with most solutions, this one has drawbacks. Java is available only on specific platforms, such as Windows 95, Windows NT, Solaris 2.x, SPARC, and Macintosh systems. No other platforms are supported by Sun as of this writing. Also, while Java is great for applications that are run over networks, if you're selling an executable, it's not a good choice. At this stage in the language's development, Java is an interpreted compiled language. You compile your code into device independent p-code (using a compiler such as javac) and run it through an interpreter, such as the java program. If you need a single executable program, you'll have to look at other alternatives.

Going Native

A list of screen library options would not be complete without considering native user-interface libraries. If you're on Windows, you can use the Win32, MFC, or OWL APIs. On UNIX, there's Motif and Open Look, etc.

Native APIs or frameworks are convenient because they're freely available with the compiler. Problems arise when you have to port to other systems. The good news is that several native APIs have already been ported to various other machines. For example, ports now exist from X-Window calls to Microsoft Windows. If you've already coded your program on one system and don't want to rewrite it, this is one option for porting your code.

Based on what I've read, and from my experience with libraries, if I were going the native route I would probably choose a cross-platform API based on MFC. I would choose MFC mainly because it supports the most platforms. Versions of the MFC API are available for UNIX, Macintosh, and even MS-DOS. MFC is also easier to work with (at least for C++ programmers) than Win32. Beyond these two things, however, I know of no compelling reason for you to abandon a framework or library you're comfortable with.

Now for the drawbacks of native APIs. Many of the ports for other systems run much, much slower than the native version. Also, while the native version is usually freely available along with the compiler, many of the other system ports require you to buy commercial products from various vendors.

If All Else Fails . . .

One last place to look for information and comparisons on portable user interface libraries is in the Platform Independent Graphical User Interface FAQ. (See Item #11). It's also helpful to keep an eye on current developments through magazines and on the Internet to see if anything new eventually surfaces.

What Works for Me

So far, I've run through several options, but none of them have worked perfectly for my situation, and perhaps would not for yours either. As I've mentioned, my biggest problem in locating a screen library is trying to find a user interface that can support both character-based and graphical systems. Many implementors of libraries (both commercial and non-commercial) feel that character-based libraries are on their way out and a waste of their time to implement or support. In answer to that thinking, I would like to quickly point out the cases when a programmer may need a character-based user


new!
PowerPC

REAL TIME MULTITASKING

Multitasking in Real and Protected Mode



smx[®] Full-featured, high-performance, preemptive kernel. Customized to x86 processors. Ideal for demanding applications. Real mode works stand-alone or with DOS. 16-bit, segmented protected mode works with pmEasy16 or 286 DOS Extender. 32-bit, flat protected mode works with pmEasy32 or TNT.



pmEasy[®] 16- or 32-bit protected mode entry, DPMI services, application loaders. Periscope/32[®] and Soft-Scope[®] debugger support.

DOS-Compatible File System



smxFile[™] Full-featured file manager. IDE, floppy, flash, RAMdisk, and PCMCIA drivers available.



smxDLM[™] Runs independent executables as tasks which may be downloaded or loaded from disk, floppy, flash, etc.

Dynamically Loadable Modules



smxNet[™] TCP/IP stack. Fast UDP. Packet driver interface. Ethernet, SLIP, and PPP drivers. FTP, SNMP, NFS, Net server, other protocols.



smxProbe[™] Provides tracing and symbolic debugging. Works with or without code debuggers. Local or remote operation.

C++ Classes



smx++[™] Class library built upon smx. Provides fully OOP-compatible kernel interface.



smxEMS[™] Allows copying data between real memory and extended memory buffers or accessing extended memory via a window.

Extended Memory, Real Mode



smxWindows[™] Text windowing. Dresses up user interfaces. **MetaWindow[®]** & **Zinc[®]** graphics support.



Royalty-free licenses. Supports common, low-cost C compilers, debuggers, and locators. Great manuals. Source code available.

User Interface



On the market 7 years. 100's of applications. Extensive error checking. 30-day free trial.



1-800-366-2491
mdi@earthlink.net

MICRO DIGITAL, INC

fax 714-891-2363 int'l 714-373-6862
<http://www.earthlink.net/~mdi>

Request Reader Service #150

Accelerated Technology, Inc.

Real-Time Solutions

that get
your application
there
on-time

Nucleus PLUS and
Nucleus RTX
Real-Time Kernels

Nucleus NET
Real-Time TCP/IP

Nucleus FILE
Real-Time MS-DOS
File System

Nucleus DBUG
and kernel-aware
debugger integration

Processor Support
68xxx • 8086/80186 • 80386/486PM
DOS • Am29xxx • i960 • IDT305x
LR330x0 • R4000 • SPARClite • ARM
PowerPC • SHx • TMS320C3x/4x/2x/5x
DSP560x0 • 68HC11/16 • M37702
H8/300H • NEC78K • VR85x

[Nucleus Real-Time Software]

call 800-468-6853 today
for more information and free demo disk!

Accelerated Technology, Inc.
Post Office Box 850245
Mobile, Alabama 36685
(334)661-5770 • fax: (334)661-5788

Request Reader Service #151

Page 60

interface library, and to remind implementors of a few reasons for using one in the first place.

One reason to prefer a character-based library is that the user may be working with a low-end machine. Believe it or not, some people out there are still running old-style 8086 or 80286 PCs with nothing more than DOS installed. Not every user is a programmer who needs to keep up-to-date with the latest systems. If you're working with limited hardware, you may not have the luxury of using a full-blown graphical operating system either. Some people actually do prefer the look of character-based programs and prefer to purchase them over other programs. It is often easier to read the print on character-based screens than on graphical screens. The moral of this story is: it's important to know who your end user is and be familiar with your market or customer base. Also, it's easier to choose a library that supports more than what you currently need, than to have to search for and port to another new library later on if you find the one you have doesn't support all your goals.

If you're in a position similar to mine, you're probably wondering if I ever found my ideal screen library. The answer is yes and no. What I finally decided to do was pick some libraries that worked well on the systems I needed to implement them on and create some wrapper classes to make those implementations all appear alike on the different systems. This solution keeps me from writing everything from scratch. For Windows I chose MFC, since code written for it can easily be ported between Windows 3.1, 95, and NT. A port of MFC is also available for the Mac platform. When I have more time, I can change the underlying code to use a native Macintosh API and increase the execution speed of my programs for that platform. For DOS platforms, I decided to use D-Flat and the D-Flat wrapper classes. The controls in this library and the D-Flat wrapper class interface are very similar to the MFC library. This helps decrease coding time for the DOS platform port.

I haven't found any one solution that's perfect for my situation; I still have had to do some coding. I created three main wrapper classes similar to the classes used in MFC. One encapsulates the application, another the frame window, and the third encapsulates the individual MDI (multiple document interface) windows. Any functionality specific to the libraries I'm working with is hidden within these three classes. I then designed a

uniform interface for the wrapper classes that lets me write to, get information from, or add controls to a window in the same way regardless of the underlying library I'm using. For example, my MDI Window class wrapper has a text function that lets me add text to the screen the same way on any platform. The member function for sending text to a window appears as follows:

```
void text (char *words,
           position *pos,
           attributes *att);
```

I also took pains to encapsulate data in classes as much as possible. This data includes information such as position of controls or text on screen, color of controls or text, and the actual text being written to or read from the screen. Not only does encapsulating this data help when porting, it also allows for room to grow. The requirements of my screen library and the way information is displayed or the type of information I display may change. By encapsulating the information, it becomes much easier to make these changes later. This also helps minimize the amount of code that will be affected by these changes.

My final screen library design is composed of two parts. I choose a platform-specific screen library for the platform I'm working on. I then encapsulate the screen library in my own classes so that I have a universal subset of commands that will get the job done. I hide the platform-specific code within these classes and use it to create my own screen library. I had to do a lot of searching to find which libraries to use, and to finally come to the conclusion that no one library is right for my specific needs. Hopefully, I've managed to give you enough concrete examples of actual user interface libraries available to make your own search for an ideal screen library easier. □

References

- [1] Al Stevens. "C Programming" column, *Dr. Dobb's Journal*, May 1991 - May 1995.
- [2] M.A. Sridhar. *Building Portable C++ Applications with YACL* (Addison-Wesley, 1996). ISBN:0-201-83276-3.
- [3] John K. Ousterhout. *Tcl and the TK Toolkit* (Addison-Wesley, 1994). ISBN: 0-201-63337-X.
- [4] Brent B. Welch. *Practical Programming in TCL and TK* (Prentice-Hall PTR, 1995). ISBN: 0-13-182007-9.



Dan Saks

Abstract Declarators, Part 3

Sometimes looking at a problem a different way makes it all come into focus. Dan presents an alternative grammar to make elements of C++ syntax much easier to understand.

Copyright © 1996 by Dan Saks

A declarator is the central part of an object or function declaration. It includes a *declarator-id* (the name being declared), along with any operators that modify that *declarator-id*. For example, the declarator in

```
unsigned char *const p[N];
```

is **const p[N]*. It declares that *p* (the *declarator-id*) has type “array with *N* elements of type *const pointer*” to something. The *decl-specifiers* (the stuff to the left of the declarator) tell us that the something is *unsigned char*.

An abstract declarator is a declarator without a *declarator-id*. Abstract declarators appear in numerous C++ constructs, probably most often in parameter declarations and cast expressions. They also appear in exception declarations, exception specifications, template argument lists, and *sizeof* expressions.

For example, a function declaration such as

```
char *memcmp(const void *,
             const void *,
             size_t);
```

Dan Saks is the president of Saks & Associates, which offers consulting and training in C++ and C. He is secretary of the ANSI and ISO C++ committees. Dan is coauthor of C++ Programming Guidelines, and codeveloper of the Plum Hall Validation Suite for C++ (both with Thomas Plum). You can reach him at 393 Leander Dr., Springfield OH, 45504-4906, by phone at (513)324-3601 (the area code changes to 937 after September 1996), or electronically at dsaks@wittenberg.edu.

has three parameter declarations, yet no parameter names. Therefore, each has an abstract declarator. In the first two parameter declarations, the abstract declarator is just ***; each parameter’s type is “pointer to *const void*”. In the third parameter declaration, the abstract declarator is actually empty; the parameter’s type is just *size_t*.

Abstract declarators appear in cast expressions as part of the *type-id* (the part of the cast that specifies the result type). For example, in the cast expression *static_cast<D &>(b)*, the *&* is an abstract declarator specifying a reference type. The complete *type-id D &* specifies the type “reference to *D*”.

For quite a few months now, I’ve been developing a program called *decl* that parses C++ declarations and translates correct ones into English. At present, *decl* recognizes declarators but not abstract declarators. In my last column (two months ago), I started augmenting the program to recognize abstract declarators, planning to finish it the following month. But, the software business being what it is, things didn’t quite work out that way. This month, I’ll just say I’ll make some more progress.

decl’s Current State

decl has two main parts: a scanner and a parser. The scanner reads characters from the input stream, filters out the whitespace characters, and partitions the remaining characters into tokens (identifiers, keywords, operators, and punctuation). The parser acquires tokens from the scanner, pieces them into declarations, rejects incorrect declarations, and translates correct ones into English.

Two months ago, I explained how recognizing abstract declarators complicates

the parser by occasionally forcing it to look ahead more than one token to determine the next parsing action. (See “C++ Theory and Practice: Abstract Declarators, Part 1,” *CUJ*, June 1996.) Last month, I gave the scanner the ability to look ahead and back up by more than one token, in preparation for adding abstract declarators to the parser. (See “C++ Theory and Practice: Abstract Declarators, Part 2,” *CUJ*, July 1996.)

It’s been several months since I presented the parser in its entirety (see “The Column That Needs a Name: Recovering from Parsing Errors,” *CUJ*, April 1996). As I mentioned, that parser recognizes declarators, but not abstract declarators. It also recognizes function declarators, but does not allow parameters in them. This month, I will explain the remaining design issues for a parser that can recognize not only abstract declarators, but also function declarations with non-empty parameter lists.

Left Recursion vs. Iteration

One of the points I’ve been trying to make with these articles on parsing is that, once you really learn to read it, the C++ grammar can guide you in synthesizing or decomposing any C++ construct. I’ve been using the *decl* program to show how the grammar maps into a fairly simple method for recognizing C++ programs, called recursive-descent parsing. Many compilers actually use recursive-descent. Understanding the technique will give you greater insight into C++ syntax and the meaning of compiler diagnostics.

Before I add new code to parse abstract declarators and function parameter lists, I want to take another look at the grammar. Table 1 shows a grammar for those C++ declarations that *decl* already recognizes. The grammar lacks some details of complete C++ declarations, but it covers a very useful subset. It also indicates that the parentheses in a *function-suffix* may surround a *parameter-list*. Until now, this *function-suffix* has been just a place-holder; *decl* does not yet recognize function parameters.

As always, I’ve expressed the grammar in the EBNF notation, summarized in Table 2. This notation is not the same as the notation used in the draft C++ Standard. Rather, the draft uses the notation that K & R (Kernighan and Ritchie [1]) used to describe C.

POWERFUL, YET AFFORDABLE

Razor from Tower Concepts, Inc. has a proven track record as a powerful, integrated software tool suite for UNIX workstations. Razor is simple to learn, easy to configure but powerful in its performance capabilities for configuration management and problem tracking. Priced at only \$495 per floating license, Razor is the best value on the market today to save development time and costs.

CONFIGURATION MANAGEMENT

The flexible Configuration Management (CM) module allows for seamless integration into your work environment. Razor users don't need to spend days immersed in manuals or training sessions on how to use Razor. Users are often up and running on Razor within 30 minutes.

Razor/CM supports ASCII or binary files, easily imports your existing SCCS/RCS investments,

is highly extensible using triggers and scripts, and requires minimal system resources. Because Razor uses a non-proprietary database, direct user queries are allowed and encouraged.

**WHAT IS
POWERFUL,
EASY TO USE,
AND A GREAT
VALUE?**

PROBLEM TRACKING

The Problem Tracking (PT) module features fully configurable forms that simplify data entry, editing and routing. And Razor's powerful

reporting capabilities provide engineers and managers with timely feedback. Razor/PT supports parallel databases at remote sites and offers an e-mail interface for information entry, query and review.

OUTSTANDING VALUE

Whether you utilize all of its features or just some, Razor's intuitive interface, matchless versatility and superb technical support makes it an excellent investment.

Available on all major UNIX platforms, Windows 95/NT and HTML interfaces. For free eval copies and full documentation, visit our web site at <http://www.tower.com>.

Razor®

103 Sylvan Way,
New Hartford, N.Y. 13413 U.S.A.
sales@tower.com
<http://www.tower.com>
315-724-3540

□ Request Reader Service #152 □



One of the problems with the K & R notation is that it has no direct notation for iteration (repetition). Grammars that use this notation must combine recursion with alternation (choices) to simulate iteration. Don't get me wrong; I like recursion, but not to the exclusion of iteration. Using recursion to describe a repetitive construct is not as clear as using iteration. And, as contradictory as this may seem, there are some recursive rules that you can't map directly into a recursive-descent parser. Permit me to explain.

Table 1: A simplified grammar for C++ object and function declarations

```

simple-declaration = decl-specifier-seq declarator-list .
decl-specifier-seq = decl-specifier { decl-specifier } .
decl-specifier = type-specifier .
type-specifier = cv-qualifier | simple-type-specifier .
simple-type-specifier = type-keyword | type-name .
type-name = name .
declarator-list = declarator { "," declarator } .
declarator = direct-declarator | ptr-operator declarator .
direct-declarator = ( declarator-id | "(" declarator ")" ) [ array-suffix | function-suffix ] .
declarator-id = identifier .
array-suffix = "[" [ constant-expression ] "]" .
constant-expression = name | integer-literal .
function-suffix = "(" [ parameter-list ] ")" cv-qualifier-seq .
ptr-operator = "&" | [ class-name "::" ] "*" cv-qualifier-seq .
class-name = name .
cv-qualifier-seq = { cv-qualifier } .

```

Table 2: A summary of EBNF operators

X = Y .	means "an X is defined to be a Y"
X Y	means "either an X or a Y"
[X]	means 0 or 1 occurrence(s) of an X"
{ X }	means "0 or more occurrences of an X"
also	
()	group sub-expressions
" "	enclose a terminal symbol

With recursive-descent parsing, each production in the grammar maps (more-or-less) into a separate parsing function. In `decl1`, the parser is a class, and each parsing function is a member of that class. For example, the grammar has a production

```
declarator =
    direct-declarator |
    ptr-operator declarator .
```

and so the parser has a member function called `declarator`. Each parsing function returns a string that spells out the meaning of what it just parsed.

The right-hand side of a typical production refers to the names of other productions. Each reference becomes a function call to the corresponding parsing function. For example, the parsing function for `declarator` includes a call to `ptr_operator` followed by a (recursive) call to `declarator`. The | (vertical bar) in a production indicates a choice. In this case, it indicates that a `declarator` can be either a `direct-declarator`, or a `ptr-operator` followed by another `declarator`.

The `declarator` parsing function appears in Listing 1. `input` is the scanner object. `input.current()` returns the current token, and `input.current().kind()` returns the current token's category. The function looks for a token that can begin a `ptr-operator` before calling `ptr_operator`. If it finds no such token, it blindly calls `direct_declarator`. I took a little liberty in translating the grammar into this function. Had I translated the grammar rigorously, the function would check for a token that can begin a `direct-declarator` before calling `direct_declarator`.

However, I wrote `declarator` assuming that `direct_declarator` will complain if it is not happy with the token it sees.

The `declarator` parsing function uses recursion, and uses it properly. `ptr-operators` bind to a `declarator` from right to left, and recursion is an effective tool for reversing the order of the `ptr-operators` so they appear from left to right in the output.

For example, the `declarator *&r` means `r` is a reference to a pointer, not a pointer to a reference. Each call to `declarator` that begins with a `ptr-operator` parses that one `ptr-operator` and saves its translation. `declarator` calls itself to parse the remainder of the `declarator` (which may include more `ptr-operators`). When the recursive call returns, `declarator` appends the translation of the `ptr-operator` to the end translation of the `declarator`. Thus, the first-in-last-out nature of recursion reverses the order of the `ptr-operators`.

Now let's look at the grammar for `direct-declarator`. The draft C++ Standard uses recursion in this rule, too:

```
direct-declarator =
    declarator-id |
    "(" declarator ")" |
    direct-declarator ( array-suffix | function-suffix ) .
```

This rule has three primary alternatives. The first two alternatives are no problem. Each begins with a unique token that leads the parser toward that choice. However, the last alternative is a puzzler. It begins with a reference to `direct-declarator`, which is this very same rule. Thus, the tokens that lead the parser to this alternative are exactly the same as the tokens that lead to the other alternatives. The

Listing 1: The declarator parsing function

```
// declarator =
//     direct-declarator | ptr-operator declarator ,
// string parser::declarator()
{
    token::category tc = input.current().kind();
    if (tc == token::AMPERSAND || tc == token::STAR
    || tc == token::NAME)
    {
        string p = ptr_operator();
        return declarator() + p;
    }
    else
        return direct_declarator();
}
//End of File
```

1.

2.

3.

You can count on the 68000 Real-Time Multitasking Kernel

Time is of the Essence. Start today with AMX 68000, KADAK's fast, full-featured, compact, ROMable kernel for Motorola MC680x0 and MC683xxx processors. And, because time to market is critical, get a head start on product construction with our Configuration Builder. Then, rapidly test with our *KwikLook™* tool which, when coupled with the popular *SDS SingleStep™* debugger, provides the ultimate in task-aware debugging.

You Can Count On KADAK. We've been setting real-time standards since 1978. Sony, Hewlett Packard, Philips, Siemens, Hughes Aircraft ... 1,000 of the most demanding companies count on AMX™, its documentation and our highly-responsive, no-nonsense approach to product support.

A Price You Can Count On. KADAK's simple is better business philosophy means you pay a fair price, once, for a site license—no royalties, no hidden charges and nothing extra for source code.

Other KADAK AMX Kernels for: protected mode 80386; 80x86/88; PowerPC™, R30xx, LR33xxx; 29K; i960®.

Contact us today for a sample copy of our *KwikLook™* debugging tool and examples of our exceptional documentation.

AMX and *KwikLook* are trademarks of KADAK Products Ltd.
All trademarked names are the property of their respective owners.

□ Request Reader Service #153 □

We thought about
getting President
Clinton to speak at
our D.C. conference.

But we figured
you'd rather hear
Tim Berners-Lee.



Register for Software Development '96 East
October 28–November 1, Washington, D.C.

As the inventor of the World Wide Web, Tim is considerably more qualified to provide you with insights into web development. As are the 30 other web authorities teaching at SD '96 East.

With 200 classes and 175 tools vendors, SD '96 East is the largest development event on the East Coast. You'll learn concrete skills from expert instructors including Bruce Eckel, Mark Pesce

and Richard Hale Shaw. Trade intelligence with the industry's smartest development professionals. And preview the latest tools for intranet development, object-oriented programming, client/server migration and rapid application development.

To get your SD '96 East conference catalog, call us at 800.441.8826 or 415.905.2702. Or send us e-mail at sd96east@mfi.com.

We'll provide you with an exceptional opportunity to meet face-to-face with some of the most influential leaders in Washington, D.C.



More tools. More intelligence.

<http://www.sd96.com>

only way the parser can distinguish the third alternative from the first two is by looking ahead for an *array-suffix* or a *function-suffix*.

Now that the scanner provides unbounded lookahead, looking ahead should be easy, but in this case it isn't. Well, maybe the looking is easy but finding the right token is hard. An *array-suffix* begins with a left bracket and a *function-suffix* begins with a left parenthesis. The parser cannot simply look for the first occurrence of one of these tokens; it must look for the last occurrence that has not been found by a previous lookahead.

For example, consider the declarator `x[10]`. The first time the parser enters `direct_declarator`, it sees `x` as the current token. It must decide if it can simply accept `x` as a *declarator-id*, or call itself and then call `array_suffix` to parse a suffix. Since there is a suffix after `x`, `direct_declarator` does call itself.

Each call to `array_suffix` parses exactly one suffix. Therefore, the recursive call to `direct_declarator` must parse just `x`, and leave the `[10]` for `array_suffix`.

Upon entering the recursive call to `direct_declarator`, the current token is still `x`. This call to `direct_declarator` can't tell whether it was called recursively or from another function such as `declarator`. It must decide for itself if it can just accept `x` as the *declarator-id*, or if it must call itself again. There is a suffix after `x`, suggesting that a recursive call is in order. However, that suffix, `[10]`, has already been spoken for (by the previous call to `direct_declarator`), so this recursive call must not parse the `[10]`. Thus, when the parser looks ahead for a suffix, it must stop before reaching the `[10]`.

Parsing a declarator such as `x[10][20]` gets even harder. In general, the parser must somehow exclude one more suffix from the lookahead on each recursive call to `direct_declarator`. This is no simple task.

The production for `direct-declarator` shown above is said to be *left-recursive*, because the recursive reference to `direct-declarator` appears as the leftmost symbol of an alternative on its right hand side. In contrast, the production for `declarator` shown earlier is *right-recursive*. Right-recursive rules translate pretty easily into recursive-descent parsing functions; left-recursive rules do not. Fortunately, you can often transform a left-recursive production into an iterative one that's much easier to parse. Let's do this for `direct-declarator`.

The left-recursion in `direct-declarator` suggests that the suffixes group from left to right. For example, in `x[10][20]`, `[20]` is the suffix for `x[10]`, and `[10]` is the suffix for `x`. Thus, `x[10][20]` is equivalent to `(x[10])[20]`. A parser can simply read the suffixes from left to right and pass their translations in that order to the output. Therefore, given `x[10][20]`, then `x` is an

"array with 10 elements of type array with 20 elements of type" something.

The first two alternatives in the production for `direct-declarator` indicate that a `direct-declarator` need not have any suffix at all. The third alternative indicates that a `direct-declarator` followed by one suffix is still a `direct-declarator`, and thus may be followed by two or more suffixes. The iterative rule for `direct-declarator` is therefore

```
direct-declarator =
  ( declarator-id | "(" declarator ")" )
    { array-suffix | function-suffix } .
```

for C/C++ Version 7.0 presents Bug # 578

```

1  #include <iostream.h>
2
3  class X
4  {
5      public:
6          X( const & X ) { kind = "copy"; }
7          X( ) { kind = "original"; }
8          char *kind;
9      };
10
11 void f( X x )
12 {
13     cout << x.kind;
14 }
15
16 int main()
17 {
18     X x;
19     f( x );
20     return 0;
21 }
```

It was the programmer's intent that whenever the class X is 'copied' through the copy constructor the value of 'kind' was to reflect that fact. But instead of printing "copy" it prints "original". Where did the programmer go wrong? Call if you need a hint, or refer to our web page at <http://www.gimpel.com>.

PC-lint for C/C++ will catch this and many other bugs. It will analyze a mixed suite of C and C++ modules to uncover bugs, glitches, quirks and inconsistencies.

Version 7 of PC-lint breaks new ground with inter-statement value tracking for both automatic variables and class data members. Taking clues from assignment statements, initializers and conditional expressions it can detect out-of-bound subscripts and potential null pointer uses. As an enabling technology, almost 100 standard functions are rigorously checked. Also macros are subject to increased scrutiny, checking for unparenthesized parameters, unparenthesized bodies and repeated arguments having side-effects.

Plus Our Traditional C/C++ Warnings: Uninitialized variables, inherited non-virtual destructors, strong type mismatches.

inadvertent name-hiding, suspicious expressions, etc., etc.

Full C++ Support - PC-lint for C/C++ is based on the ARM and is tracking the latest ANSI/ISO draft including exceptions and templates. It supports both Borland and Microsoft C/C++.

PC-lint for C/C++ \$239
 Numerous compilers/ libraries supported.
 Runs on MS-DOS (Optional built-in 386 DOS extender), OS/2, NT and Windows 95.
 This price is subject to increase.

FlexeLint for C/C++
 The same great product for other operating systems. Runs on all Unix systems, VMS, mainframes, etc. Distributed in shrouded C source form. Call for pricing.

Gimpel Software
 3207 Hogarth Lane, Collegeville, PA 19426
CALL TODAY (610) 584-4261 Or FAX (610) 584-4266
 30 Day Money-back Guarantee.
PA add 6% sales tax.

PC-lint and FlexeLint are trademarks of Gimpel Software

This is the way it appears in Table 1. The parsing function for this production requires no lookahead at all.

Syntax vs. Semantics

Now let's look at the grammar rules for abstract declarators and parameter lists. A *parameter-list* is a comma-separated list of parameter declarations. The draft C++ Standard describes *parameter-list* using a left-recursive production:

```
parameter-list =
  [ parameter-list "," ] parameter-declaration .
```

Table 3: A grammar for parameter-declaration using left-recursion

```
parameter-declaration =
  decl-specifier-seq ( declarator | [ abstract-declarator ] ) .

declarator =
  direct-declarator |
  ptr-operator declarator .

direct-declarator =
  declarator-id |
  "(" declarator ")" |
  direct-declarator ( array-suffix | function-suffix ) .

abstract-declarator =
  direct-abstract-declarator |
  ptr-operator [ abstract-declarator ] .

direct-abstract-declarator =
  "(" abstract-declarator ")" |
  [ direct-abstract-declarator ] ( array-suffix | function-suffix ) .
```

10450 Stancliff, Suite 110
Houston, TX 77099-4383

FAX: 713-561-9980
Phone: 800-525-4302 or 713-561-9990
e-mail: sales@esphou.com

□ Request Reader Service #154 □

Rewriting this using iteration yields a rule that's easier on the eyes as well as the parser:

```
parameter-list =
  parameter-declaration { "," parameter-declaration } .
```

The draft describes a *parameter-declaration* as:

```
parameter-declaration =
  decl-specifier-seq
    ( declarator | [ abstract-declarator ] ) .
```

This is similar to a *simple-declaration*, except that (1) a *parameter-declaration* can have at most one *declarator* (rather than a comma-separated list of *declarators*), and (2) the *declarator* in a *parameter-declaration* can be abstract.

Remember, an abstract declarator is just like an ordinary declarator, except it doesn't have a *declarator-id*. Thus, the productions for *declarator* and *abstract-declarator* should be nearly identical. Yet the draft C++ Standard opts to keep them as distinct productions, rather than combine them into one. This approach has both advantages and disadvantages.

There are places in the grammar, such as *simple-declaration*, that allow declarators but not abstract declarators. There is also one place, *type-id*, which allows abstract declarators but not ordinary (non-abstract) declarators. If you combine the productions for abstract declarators with declarators, it seems that you lose the ability to distinguish the places which allow only one kind of declarator.

On the other hand, ordinary declarators and abstract declarators are nearly identical, and describing them with different productions obscures their similarities. Table 3 shows the productions for *declarator* and *abstract-declarator* in the left-recursive form used in draft C++ Standard. As expected, the productions for *declarator* and *direct-declarator* are similar to the productions for *abstract-declarator* and *direct-abstract-declarator*, respectively. (The similarities become more apparent if you delete the word *abstract* everywhere.) However, considering that the only difference between a *declarator* and an *abstract-declarator* is the presence or absence of a *declarator-id*, it's surprising that the productions don't look more alike.

The other problem with using different rules for the different kinds of declarators is that it introduces parsing problems that require elaborate lookahead. Merging *declarator* and *abstract-declarator* into a single set of productions solves most of those problems.

As a first step in merging the productions, I removed the left-recursion and replaced it with iteration. The resulting productions appear in Table 4. I believe these productions make the similarities between *declarators* and *abstract-declarators* more evident. In fact, the only difference (other than the presence of *abstract-* in a bunch of places) is that the *declarator-id* is missing from a *direct-abstract-declarator*.

The productions in Table 4 still pose parsing problems that require looking ahead. A *direct-abstract-declarator* can begin with a left parenthesis, followed by an *abstract-declarator* and a right parenthesis. The parenthesized *abstract-declarator* is optional, so the parser can skip it and go on to look for an *array-suffix* or a *function-suffix*. However, a *function-suffix* also begins with a left parenthesis. Therefore, the parser must look at the next two tokens to distinguish a *function-suffix* from a parenthesized *abstract-declarator*.

For example, consider the abstract declarators (`T &`) and (`T::*`). The first declarator, (`T &`), is a *function-suffix*. The second, (`T::*`), is a parenthesized *abstract-declarator* describing a pointer-to-member. In parsing either of these declarators, the parser enters `direct_abstract_declarator` with the left parenthesis as the current token. If the parser looks at the next token after the parenthesis and sees a type name such as `T`, it still can't tell if it has a parenthesized *abstract-declarator* or a *function-suffix*. It must look at one more token. If that token is not a `::`, then the parenthesis is the start of a *function-suffix*; otherwise, it's grouping around an *abstract-declarator*.

The lookahead problem I just described arises regardless of whether the grammar uses left-recursion or iteration. However, the iterative grammar in Table 4 has one problem that the left-recursive grammar in Table 3 does not. According to Table 4, an *abstract-declarator* can be empty. According to Table 1, the *parameter-list* (between the parentheses) in a *function-suffix* can be omitted. Therefore, a *direct-abstract-declarator* that is just `()` could be either parentheses around an empty *abstract-declarator*, or a *function-suffix* with the *parameter-list* omitted. This is a *syntactic ambiguity*. That is, there are two different and apparently valid ways to parse the same sequence of tokens.

A careful (very careful) reading of the left-recursive grammar in Table 3 shows that it does not share this problem. According to Table 3, an *abstract-declarator* cannot be empty. Therefore, a *direct-abstract-declarator* that's just `()` must be a *function-suffix*.

I prefer the grammar in Table 4. It's more readable and easier to map into a parser. The syntactic ambiguity it introduces is most unfortunate. Does this mean we must abandon it? I don't think so. Such ambiguities occur elsewhere in C++, and even in C. Both languages resolve the ambiguities by simply imposing additional semantic rules.

For example, according to C++ grammar, a function declaration such as

```
int f(const T);
```

has either an unnamed parameter of type `const T`, or a named parameter `T` of type `const` (and by default) `int`. The grammar is ambiguous. C++ resolves the conflict according to a semantic rule called the "maximal munch" rule. (See "The C++ Column That Needs a Name: Understanding C++ Declarations," *CUJ*, December 1995.) According to maximal munch, if `T` is a type name in the scope enclosing the function declaration, then the function has an unnamed parameter of type `const T`. Otherwise, it has a named parameter `T` of type `const int`.

The C++ standards committee recently discovered a syntactic ambiguity in the construct `~X()`. According to the grammar, `~X()` means either

- `~(X())`, which applies the `~` operator to a temporary object of type `X` constructed with no arguments, or
- `this->~X()`, which calls the destructor for type `X` applied to `*this`.

The committee opted to add a semantic rule favoring `~(X())`.

I usually prefer an unambiguous grammar to an ambiguous one that resolves ambiguities by semantic rules, but I'm more concerned about the overall utility of the language description. I want the production for *direct-abstract-declarator* in iterative form, so I can combine it with *direct-declarator*, and I can't use the iterative rule

for *direct-abstract-declarator* without some rule to specify the meaning of `()`. That rule is: a *direct-abstract-declarator* consisting of just `()` is always a *function-suffix*.

Now, at last, I can merge *abstract-declarator* with *declarator* to form a single set of productions. Table 5 shows the resulting simplified productions for *parameter-declaration* and *declarator*.

Merging abstract declarators with declarators introduces a problem with *declarator-list* (used in *simple-declaration*), which can accept declarators but not abstract declarators. The solution to the problem is a couple more semantic rules. Add the following production to the grammar:

Table 4: A grammar for parameter-declaration, where direct-declarator and abstract-direct-declarator use iteration

```
parameter-declaration =
    decl-specifier-seq ( declarator | abstract-declarator ) .
```

```
declarator =
    direct-declarator | ptr-operator declarator .
```

```
direct-declarator =
    ( declarator-id | "(" declarator ")" )
    [ array-suffix | function-suffix ] .
```

```
abstract-declarator =
    direct-abstract-declarator | ptr-operator abstract-declarator .
```

```
direct-abstract-declarator =
    [ "(" abstract-declarator ")" ]
    [ array-suffix | function-suffix ] .
```

Attention: Application Developers

Tom Sawyer Software's Graph Layout Toolkit will save you years of back-end development work!

- The Graph Layout Toolkit provides a portable graph data model and suite of layout algorithms that can be easily integrated into your applications. These logical layout tools calculate the (x,y) coordinates for all of your objects, even for the most complicated graphs.
- Orthogonal layout is ideal for the display of physical and logical database design diagrams, software reengineering applications, or CASE applications. It features horizontal and vertical line routing for great readability.
- Incremental hierarchical layout is ideal for workflow, project scheduling, reverse engineering, visual programming, compiler, and information management applications.
- Circular and symmetric layout are ideal for network visualization applications.
- Integrated navigation manages information that spans many graphs. Interactively show and hide the detail behind your nodes.
- We provide C++ class libraries that include ANSI C interfaces. They may be used with any windowing toolkit. Microsoft Windows (3.1, NT, 95), OS/2, Macintosh, and many UNIX variants are supported.
- Enhance the visual appeal of your applications by licensing our supported graph layout technology.

Tel: (510) 848-0853, Fax: (510) 848-0854,
E-mail: info@TomSawyer.COM
<http://www.TomSawyer.COM>



TOM SAWYER SOFTWARE
1828 FOURTH STREET
BERKELEY, CALIFORNIA 94710



Table 5: A grammar for parameter-declaration, with declarator and abstract-declarator as one set of productions

parameter-declaration = decl-specifier-seq declarator .

declarator = direct-declarator | ptr-operator declarator .

direct-declarator = [declarator-id | "(" declarator ")"] { array-suffix } function-suffix)

concrete-declarator = declarator .

along with a rule explaining that a *concrete-declarator* must have a *declarator-id*. Also add:

abstract-declarator = declarator .

along with a rule explaining that an *abstract-declarator* must not have a *declarator-id*.

Rules such as these already exist in both C++ and C. For example, they both have the production:

constant-expression = conditional-expression .

along with a slew of rules constraining the operands and operators that can appear in a *constant-expression* (so the compiler can evaluate the expression at compile time). The alternative is a lot of extra productions to describe *constant-expression*.

With these productions that distinguish *abstract-declarators* from *concrete-declarators*, the production for *declarator-list* becomes:

declarator-list = concrete-declarator { "," concrete-declarator } .

I'll show you the code for the revised parser next time. For those of you who can't wait, you can find it on *CUJ's* ftp site (see page 3 for ftp information).

Errata

In past articles, I presented versions of the *decl* grammar that defined *decl-specifier-seq* as:

decl-specifier-seq = { decl-specifier } .

This suggests that the sequence can be empty. It cannot. The grammar in Listing 1 correctly shows that the sequence must contain at least one *decl-specifier*.

Thanks to Steven Lee (slee@Cisco.com) for bringing this to my attention. □

Reference

- [1] Brian Kernighan and Dennis Ritchie. *The C Programming Language* (Prentice Hall, 1978).

Learn C and C++ Faster with Expert Video Instruction

Two comprehensive video courses teach you how to program in C and C++ in the shortest time possible. Developed by training specialist Silicon River Ltd., the video courses result from years of experience in teaching programmers how to take advantage of these powerful languages.

The video courses are very well structured, presenting the material in a logical progression that makes it easier to comprehend. Advanced display techniques are used throughout to make the discussion of language features interesting and clear.

The languages are covered in detail. Each course includes six video tapes (approx. 13 hours total length), a workbook, and a disk full of example programs. The courses are divided into units that are studied one at a time. For each unit, you first view the video presentation, then work through the corresponding exercises in the companion workbook. The programming exercises are designed to re-enforce the concepts presented in the video. After completing the exercises, you return to the video to view a brief workshop segment. The workshop is an interaction between instructor and students, reviewing the assignments just completed. This allows you to compare your experiences with those of other

students taking the course. Including the time spent working through the exercises, most students are able to complete each course in approximately 60 hours.

```
#include <stdio.h>
#include <stdlib.h>
#include "video.h"

main()
{
    FILE *fp, *tp;
    int Input = 0, Output;
    char *FileName =
        malloc(L_tmpnam);
    if (!FileName)
        Error(V_NO_MEMORY);
```

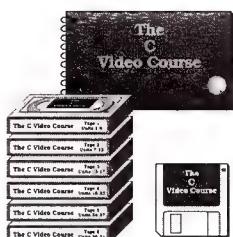


The C Video Course covers the Standard C language. Since C++ is an extension of the C language, it is recommended that you learn C before attempting to learn C++. The C Video Course prepares you for the C++ Video Course, which focuses exclusively on C++ specific language features and object-oriented concepts.

The C and C++ video courses are suitable for either home or office use. All you need is a TV and VCR to view the tapes. Of course you will need a computer and a C or C++ compiler to perform the exercises. Since both courses focus on standard language features, any Standard C or C++ compiler may be used.

By themselves, these two video courses are very effective at teaching you C and C++. But you can also use them in conjunction with live training that you might receive at a seminar or a local college. The videos provide a convenient way to review subjects that you may not completely understand. Companies interested in using the video courses to train several programmers at once should ask about cost-effective corporate packs.

□ The C Video Course



List Price: \$299.90

□ The C++ Video Course



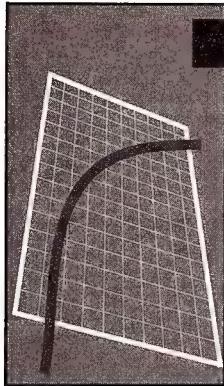
List Price: \$299.90

To Order Call:
1-800-333-0330
or
Tel: 1-214-783-6001
Fax: 1-214-783-1404

Mix Software
1132 Commerce Drive
Richardson, TX 75081

<http://www.mixsoftware.com>

□ Request Reader Service #156 □



Bobby Schmidt

C->C++ Mutations, Part 4

In this installment, Bobby wraps up his series on porting C code to C++. Once again he shows that adapting to C++'s stricter rules is painful, but ultimately results in better code.

Copyright © 1996 Robert H. Schmidt

What you have in your hot little hands is the fourth and final exploration of C code that mutates behavior when ported to C++.

Type of Character Literals

In the Cretaceous era of Classic C (K&R C), everything seemed to be an int. Unspecified function return and object types were ints. Pointers were ints. shorts were ints (and often still are). longs were ints. floats and doubles weren't ints, but then, real programmers didn't use floating-point [1]. And while chars typically weren't ints, char arguments were promoted to ints, and lots of programmers mixed int and char like beer and pretzels. After all, who among us has never seen or written constructs like [2]

```
char index_to_letter(int const index)
{
    return 'A' + index;
}
```

Standard C has maintained much of the int legacy. Given this trend, you should be nonplussed to learn that so-called "character" literals like 'A' are actually of type int. This leads to the interesting anomaly that, for an ASCII execution character set,

```
char c = 'A';
```

Bobby Schmidt is a freelance writer, teacher, consultant, and programmer. He is also a member of the ANSI/ISO C standards committee, an alumnus of Microsoft, and an original "associate" of (Dan) Saks & Associates. In other career incarnations, Bobby has been a pool hall operator, radio DJ, private investigator, and astronomer. You may summon him at 3543 167th Ct NE #BB-301, Redmond WA 98052; by phone at (206) 881-6990, or via Internet e-mail as rschmidt@netcom.com.

is morally equivalent to

```
char c = 65;
```

Even so, I encourage you to use 'A' here. 'A' is much more self-documenting, implying that you are thinking of the logical abstracted character 'A', and not the physical numerical implementation 65. Also, your code will port correctly to a non-ASCII environment, since 'A' will map to the appropriate encoding for the new target.

That 'A' is an int is typically invisible to you. The one place you may become aware of it is in statements like

```
if (sizeof('A') == sizeof(int))
    printf(
        "'A' is the size of an int\n");
else if (sizeof('A') == sizeof(char))
    printf(
        "'A' is the size of a char\n");
else
    printf(
        "'A' is of mysterious size\n");
```

This code, built as C, produces

'A' is the size of an int

The same code built as C++ produces

'A' is the size of a char

The reason, of course, is that character literals in C++ are truly of type char. Why this change from C? To better support function overloading, a concept foreign to C. Consider the two C++ overloaded functions

```
void spew(char const value)
{
    printf(
        "the value is '%c'\n",
        (char) value);
```

```
}
void spew(int const value)
{
    printf("the value is %d\n",
        (int) value);
}
```

If you call spew as

```
spew('A');
```

what do you expect to see? Were 'A' still of type int, you'd get the unintuitive result

the value is 65

Because C++ reckons 'A' as a char, you actually get the more useful

the value is 'A'

Where Types are Declared

You may think, rationally enough, that all C types must be declared in declarations. But C lets you declare types in expressions also, as demonstrated by the C snippet

```
void *p = (struct s {char c;}) * 0;
struct s x;
```

Here the new type struct s is declared in the type-cast expression, allowing us to define objects (like x) of that struct type. C++ disallows such type-declaring expressions. In C++, all types must be declared in declarations.

One possible C++ rewrite of the above is

```
struct s
{
    char c;
};
```

```
void *p = 0;
s x;
```

which not only changes where the type is declared, but also takes advantage of the tagged name s appearing in the untagged name space.

Vaulting Past Object Initializations

The apparently simple definition

```
purple_dofus Barney = 1;
```

has very different meanings, depending on source language. In C the statement (loosely) means

- Reserve enough storage for a purple_dooft-sized object.

- Name that storage Barney.

- Set the initial value of that storage to 1.

In C++ the definition (again loosely) means

- Reserve enough storage for a purple_dooft-sized object.

- Name that storage Barney.

- Call the purple_dooft constructor that accepts a single integer argument, passing 1 as that argument, and passing &Barney as the constructor's hidden this pointer.

Put another way, initialization in C sets a bit pattern, while initialization in C++ calls a constructor function. That C++ constructor may generate side effects (e.g., allocating resources) requiring collateral side effects (freeing resources) in the destructor. Were an object not properly initialized, use and destruction of that object could produce disaster.

Since C initialization calls no construction function, there are no such side-effect concerns. At worst, an uninitialized C object simply contains the wrong bit pattern, a problem easily correctable by assignment. This laxity allows C compilers to accept code like

```
void f(int const i)
{
    goto jail;
    if (i)
    {
        int const x = 1;
```

High Performance Raster Imaging

RasterMaster™ and SnowVue™ - acclaimed raster imaging tools and viewers - over 100 functions. We are the developer's choice for speed, capability, features and reliability. Our TIFF and JPEG beat everyone else!

Functions

Reading & Saving
Compression
Display, rotate, zoom,
pan; scroll, & more
Auto Deskew
Despeckle
Anti aliased display
Auto aspect ratio
Low level functions
Color reduction
Image processing
Twain scan support
Internet GIF support
Large image support

Formats

40+ including G4, G3,
all TIFF, JPEG, IOCA,
CALS, ABIC, PNG,
PCX, BMP, GIF, PCD,
TGA, ASCII, EPS,...

Platforms

Windows 3.1, 3.11
Windows 95/NT
Macintosh, UNIX ...

Languages

DLL's, VBX's, OCX
Best API in the industry

Call now for free, no risk evaluations!

See our Website: <http://www.tiac.net/users/simwiz>

Snowbound Software PO Box 520 Newton, MA 02159

Tel: 617 926-1822 Fax: 617 926-1233

Imaging Software - Powerful, Fast, Reliable

□ Request Reader Service #157 □

```
jail:
    printf("x = %d\n", (int) x);
}
}
```

Because the goto skips x's initialization, you have no clue what will get written by the printf statement.

C++ expressly disallows such leaps past object initialization. Allowing them would induce run-time checks at object destruction, to see if the object's constructor had indeed been called, and would almost certainly lead to unpleasant object behavior.

No Returned Value From Function

Some C functions don't want to return anything useful. In Standard C these functions are prototyped to return void. Cretaceous era C didn't have void, so such functions were always declared to return something. Although their compile-time declarations suggested they had to return a value, at run time these functions didn't actually have to. This prevented programmers from needing to return a dummy value just to satisfy the compiler, when they knew the value would be unexpected and ignored by the caller.

Standard C retains this feature, allowing

```
int initialize(char *p)
{
    *p = '\0';
}
```

C++ has had void from the get-go; since there's no compelling reason to write functions like initialize, C++ does not allow this.

Storage Class Specifiers

C supports two storage class specifiers, extern and static [3]. The very phrase "storage class specifier" suggests that these keywords affect things occupying run-time storage. While objects of a particular type take up such storage, the types themselves do not. Nonetheless, according to the C grammar, storage class specifiers are valid (and ignored) in type definitions, giving rise to

```
static struct s
{
    int i;
};
```

The C++ grammar prevents such constructs. Among other things, preventing these constructs avoids the confusion that could arise between declaring a class member static, and declaring the entire class static.

Note that, while extern types are disallowed in C++, some C++ implementations hold the notion of exported types. For example, Microsoft Windows DLLs export entities they want to be visible to other applications or DLLs. Traditionally those entities have been objects and functions; however, with the rise of complex class libraries like MFC, DLLs also expose classes, which are conceptually compile-time entities. Exposing a class enables a system to share one copy of the library among all applications, instead of duplicating the class library object code in every client application. Such type exporting is non-standard and definitely non-portable across platforms.

Uninitialized const Objects

As bizarre as it may seem, C lets you define uninitialized const objects. That is, you can create an object in C with random values that you can never change, as in

```
double const pi; /* ooops, forgot to initialize */
```

```
double degrees(double const radians)
{
    return radians * 180 / pi;
}
```

C++, with its stronger object integrity, requires you to initialize a const object. That C does not require this is quite unfortunate. I don't know why C has this weakness; I'm guessing it's to help accommodate Classic C code ported to Standard C (as Classic C did not have the const keyword).

Type of enum

Classic C also didn't have enums. The quintessential work-around was a series of macros, such as

```
typedef int Roy_G_Biv;

#define RED 1
#define ORANGE (RED + 1)
#define YELLOW (ORANGE + 1)
#define GREEN (YELLOW + 1)
#define BLUE (GREEN + 1)
#define INDIGO (BLUE + 1)
#define VIOLET (INDIGO + 1)
```

Ick. As I explored in my January article (first of the Boolean series), macros are evil. Standard C enums are far superior:

```
enum Roy_G_Biv
{
    RED = 1,
    ORANGE,
    YELLOW,
    GREEN,
    BLUE,
    INDIGO,
    VIOLET
};
```

Unfortunately, Standard C leaves a gaping type hole in its enum support. In C, enumerators (constants like RED, defined in an enum definition) are of type int, meaning enumeration objects can be initialized and assigned from int:

```
enum Roy_G_Biv chroma = INDIGO; /* OK */

enum Roy_G_Biv colour = 6; /* also OK; identical
                           to above example */
```

continued on page 75

Better Productivity Tools!

TE Developer's Kit: Rich Text Edit Control

Incorporate advanced text editing features into your application easily and cost effectively.

Basic features: Wysiwyg multiple fonts, point sizes and character styles; paragraph indentation, centering, justification, double spacing, borders and shading; left, right, center and decimal tabs and ruler; RTF import/export, imbed picture and OLE objects, read-only mode, cut/paste, printing, and search/replace.

Advanced features: Pagination, multiple column and multiple sections, hyperlink support, hidden and protected text, tables, page header and footer, text/picture frames, embed controls, line/boxes, print preview. Includes DLL, VBX, OCX, and MFC class interface. Also includes the complete 'C' source code. (*Windows or WIN32: \$389*).

DOS and OS2 versions also available.

Also available an HTML viewer control add-on for TE (\$299).

ReportEase Plus: Report Writer Engine

ReportEase Plus consists of a report layout editor and a report executer. Advanced features include: multiple files, multiple sorts; data, calculation, system, and dialog fields; bold, underline, italic formats; subtotals, record filter, functions, word wrapping and more. This DLL/VBX features a graphic form editor with line/box drawing, colors/shades, drag/drop, print preview, etc. Simple interface works with any application. Includes the 'C' source code. (*Windows: \$459, WIN32: \$459*)

Also available a DOS version. (*\$389*)

Spell Time

Spell Time consists of a dictionary (over 100K words) and the routines to access the dictionary. It also includes an application specific dictionary, and a user dictionary. The routine will suggest alternatives for a misspelled word. Highly optimized: 400 to 500 words/sec on a 33 Mhz 386 computer. An interface with TE included. Includes the complete 'C' source code. Specify DOS, Windows, WIN32 or OS2-PM. (*\$389*)

Rich Text Table

This editable control allows the table cells to have text with multiple fonts, styles, and paragraph attributes. Other features include row and column spanning, embedded pictures, OLE, cell text scrolling and cell protection. Includes the complete 'C' source code. (*Windows or WIN32: \$399*).

ChartPro

This DLL draws bar, pie, line, area, xyz point chart, and hilo presentation graphs in unlimited styles. Features 3d rotation and dialog boxes to edit chart parameters. Includes the 'C' source code. (*Windows: \$399, WIN32: \$399*).

Sub Systems, Inc.

1-800-447-6819

Fax: 508-352-9019

11 Tiger Row, Georgetown, MA 01833, 508-352-9020

□ Request Reader Service #158 □

Advertiser Index

You can now get additional information about products and services you see advertised three ways!



cujrs@mfi.com

- ① Contact the vendor directly using the information in the advertisement. Tell them you saw their ad in *C/C++ Users Journal*.
- ② Fill in the Reader Service number found under the ad or in this index on the bound-in Reader Service Response Card. Provide the requested contact information, then mail or fax (913-841-2624) the card back to us.
- ③ Email your request to cujrs@mfi.com. Make sure you include the magazine's issue number (14.09) and where you would like the information sent.

Advertiser	Reader Service #	Page	Advertiser	Reader Service #	Page
Accelerated Technology, Inc.	151	60	D&L Online Inc	173	101
AccuSoft Corp	103	1	DC Micro Development	174	103
AI Stevens Cram Course on C/C++	**	56	Diamond Software GmbH	163	80
Aladdin Software Security	126	28	Dundas Software Ltd	107	7
Aladdin Software Security	135	40	Dyad Software	175	98
American Segment Conversion	169	100	Embedded System Products	154	66
Az-Tech Software	**	101	EMS Professional Shareware	176	101
BBT	147	54	Eonic Systems NV	162	79
Black Ice Software	159	74	FairCom Corp.	102	C4
Blinkinc	142	49	General Software	132	37
Blue Sky Software	136	42	Geodesic Systems, Inc	134	39
Borland	108	9	Gimpel Software	**	65
C Associates	170	99	Gimpel Software	**	76
C/C++ Users Journal CD-ROM	**	34	Greenleaf Software	100	C2
C++ and Windows Programming	**	47	GSK GmbH	148	55
C++ World	**	41	Haley Enterprise	177	102
CAD-UL	131	36	I-MODE Publications, Inc	129	32
Catenary Systems	137	44	Imagix	178	103
Chirp Technical Services	171	102	Information Modes	179	99
Cobalt Blue	139	46	InfoSphere	180	103
Computerpoint	172	101	Just Logic	133	38
Crystal Services	**	18	KADAK Products	153	63

Advertiser	Reader Service #	Page	Advertiser	Reader Service #	Page
KL Group, Inc.	141	48	Sanders-Indev, Inc	146	53
LEAD Technologies	124	27	Sandstone Technology	106	6
LexSaurus Software, Inc.	181	103	Scientific Placement, Inc.	196	101
Machine Independent Software	140	46	SD '96 East	**	64
Manpower Technical Services	182	98	Sequiter Software	101	C3
MarshallSoft Computing, Inc.	183	100	Sheet Bend Software	197	101
Mesquite Software	145	53	Sigma Software, Inc.	198	103
Micro Digital Inc.	150	59	Smoky Mountain Technology	127	30
Micro-Processor Services	184	100	Snowbound Software	157	70
MicroGOLD SOFTWARE	185	102	Softel vdm	**	75
MicroGOLD SOFTWARE	144	52	Software Blacksmiths	199	100
MIX Software	156	68	Software Blacksmiths	138	45
MIX Software	168	87	Software Security, Inc.	120	23
Mortice Kern Systems, Inc.	105	5	Stingray Software	160	37
National Engineering Resources	186	103	Sub Systems, Inc.	158	71
NobleNet	117	20	Systems Support Group	115	16
Nombas, Inc.	123	26	Syware	161	78
NuMega Technologies	110	11	Syware	200	98
ObjectSpace	122	25	T&T Computer Products	166	83
Omega Point Inc.	111	12	Tom Sawyer Software	155	67
On Time Marketing	143	50	Tower Concepts	152	62
Opt-Tech Data Processing	187	99	Traveltech Staffing, Inc	**	99
Parasoft	119	22	Trio Systems	164	81
Parsifal Software	188	99	ULTRA Financial Systems	201	102
Peripherie	189	99	UnderWare	118	21
Phar Lap Software Inc.	104	4	V Communications Inc	109	10
Powersoft	114	15	Vireo Software	113	14
Quadron	130	33	Voice Information Systems, Inc	202	99
R&D Books	165	82	Volt Computer	203	101
R-Active	190	103	Walnut Creek CDROM	167	85
Raima Corp	208	104	Web '96 East	**	91
Rainbow Technologies	116	17	Western Wares	204	103
Reo Centered Publishing	191	98	WIBU-Systems AG	125	27
Rincon Research Corp	192	101	Willies' Computer Software Co	128	31
Ristanovic Case	149	58	Windbase Software	205	99
Ritz Systems	193	98	Wishbone Publishing	206	101
Rogue Wave Software	112	13	Wonderware	121	24
RSVP Services	194	100	WROX Press	207	102
Ryle Design	195	99			

This index is provided as a service to our readers. The publisher assumes no liability for errors or omissions.

***This advertiser prefers to be contacted directly.*

The **TOTAL** **Fax imaging** **SOLUTION**

- ▶ Fax Development **Toolkit**
- ▶ Print Drivers for Windows, Win 95 or NT
- ▶ **Color Faxing**
- ▶ **TIFF SDK** *compression*
- ▶ **Fax C++** FOR WINDOWS, Win 95 or NT
- ▶ Internet Fax Server
- ▶ **OCX Custom Controls**

Now there's an easier and faster way to build FAX/Voice/Imaging applications. Black Ice Software provides the tools to build B/W and color applications within a few days. Free of any royalties and multi-platform compatible.

- ▶ *Image* **SDK PLUS**
- ▶ **Royalty FREE**
- ▶ **GammaLink Fax Boards**
- ▶ **Brooktrout Fax Boards**
- ▶ **Pure Data Fax Boards**

Call today for more details!

BLACK ICE
SOFTWARE INC.

292 Route 101 Amherst, NH 03031

Tel: (603) 673-1019 Fax: (603) 672-4112
BBS: (603) 673-6617 Email: blackice@mv.mv.com
CompuServe: 72662,3311

continued from page 71

As I'm sure you've guessed, C++'s type checking does not allow this. Enumerators are of their defining enum's type, meaning enum objects can be initialized and assigned only from the same enum type. If you are porting C code, you can throw in a type cast like

```
Roy_G_Biv colour = static_cast<Roy_G_Biv>(6);
```

to get around C++'s stronger rules. I find this to be a hack. Enumerations are generally supposed to be abstractions. If you find yourself relying on, or even being aware of, the underlying encoding for an enumerator, you may want to revisit your design.

Another way your ported code may break is if you assume the size of an enum. In C,

```
sizeof(INDIGO) == sizeof(int)
```

since enumerators are ints. In C++

```
sizeof(INDIGO) == sizeof(Roy_G_Biv)
```

since enumerators are of their defining enum's type. The size of an enum is implementation defined, and it may well be that ints and enums are the same on your compiler. This is not guaranteed, however, so in general portable code should not make this assumption.

Initializing char arrays

This next one helps fix the sort of bugs that drive programmers to become yak herders. Consider my self-aggrandizing program

```
#include <stdio.h>

typedef char const name_part[7];

struct name
{
    name_part last,
    name_part first;
};

struct name moi = {"Schmidt", "Robert"};

int main(void)
{
    printf("CUJ's coolest dude is %s %s!",
          (char *) moi.first,
          (char *) moi.last);
    return 0;
}
```

This code compiles and links as C, but when run on my Mac gives the less than desirable result

CUJ's coolest dude is Robert SchmidtRobert!

Who is that, the second cousin of Bond JamesBond? No, it's the result of a subtle bug that snuck by C. C++ is smarter [4], refusing to compile the code as written. If I change the definition of part_name to

```
typedef char part_name[8];
```

the code builds cleanly in both languages, yielding the much more reasonable and eminently accurate

CUJ's coolest dude is Robert Schmidt!

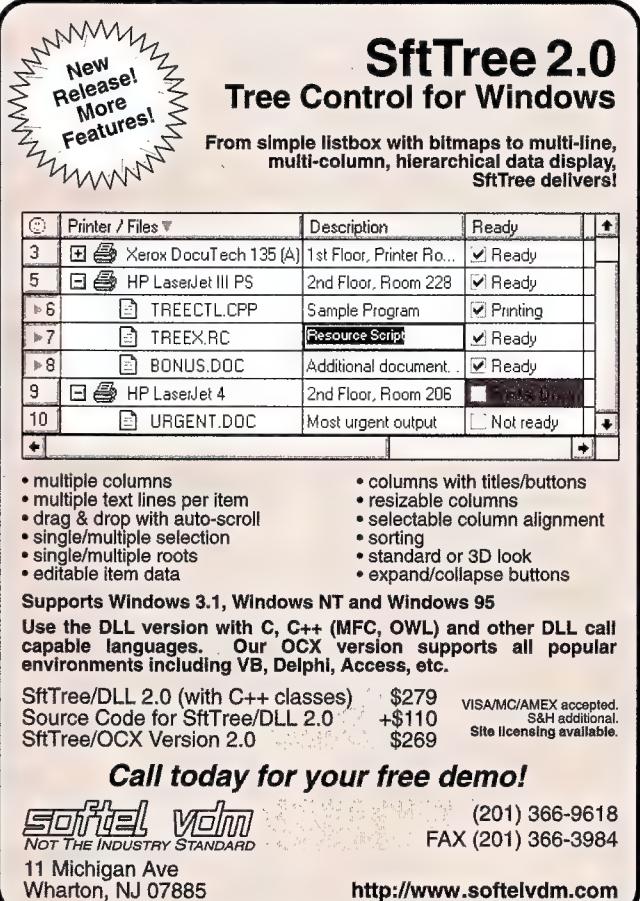
The problem was that the last array was too short to hold its initialization value: last was seven characters long, while "Schmidt" is eight (the seven letters plus the terminating '\0' byte). The letters made it on board, but that final '\0' byte got left at the altar. Because printf didn't find the terminating byte in last, it sailed past the end of that array, marching right into the immediately following storage (first).

C++'s semantics require that a char array be defined long enough to contain its initializer. The first version fails that requirement, leading to the compiler diagnostic.

Note that your results for the incorrect C version may differ from mine. For example, a compiler could align such objects on four-byte boundaries, so that up to three extra bytes of padding existed between last and first.

It's Miller (Freeman) Time!

These past four months, I've highlighted all the big changes, and most of the little ones, that you'll encounter porting valid



SftTree 2.0
Tree Control for Windows

New Release!
More Features!

From simple listbox with bitmaps to multi-line, multi-column, hierarchical data display, SftTree delivers!

Icon	Printer / Files	Description	Ready
3	Xerox DocuTech 135 (A)	1st Floor, Printer Room	<input checked="" type="checkbox"/> Ready
5	HP LaserJet III PS	2nd Floor, Room 228	<input checked="" type="checkbox"/> Ready
6	TREECTL.CPP	Sample Program	<input checked="" type="checkbox"/> Printing
7	TREEX.RC	Resource Script	<input checked="" type="checkbox"/> Ready
8	BONUS.DOC	Additional document	<input checked="" type="checkbox"/> Ready
9	HP LaserJet 4	2nd Floor, Room 206	<input type="checkbox"/> Not ready
10	URGENT.DOC	Most urgent output	<input type="checkbox"/> Not ready

- multiple columns
- multiple text lines per item
- drag & drop with auto-scroll
- single/multiple selection
- single/multiple roots
- editable item data
- columns with titles/buttons
- resizable columns
- selectable column alignment
- sorting
- standard or 3D look
- expand/collapse buttons

Supports Windows 3.1, Windows NT and Windows 95
 Use the DLL version with C, C++ (MFC, OWL) and other DLL call capable languages. Our OCX version supports all popular environments including VB, Delphi, Access, etc.

SftTree/DLL 2.0 (with C++ classes) \$279 Source Code for SftTree/DLL 2.0 +\$110 SftTree/OCX Version 2.0 \$269

VISA/MC/AMEX accepted.
 S&H additional.
 Site licensing available.

Call today for your free demo!

softel vdm
NOT THE INDUSTRY STANDARD

(201) 366-9618
 FAX (201) 366-3984
<http://www.softelvdm.com>

Standard C code to C++. Oh, to be sure, there are a few other minor examples; but they are so turgid and dull, you would risk slipping into a coma upon their mention. The full canon of such mutations appears as Appendix C of the Standard C++ Working Paper, and in fact I've based some of my examples on those given in that Appendix. By the time you read this, the committees may have foisted a new version of the Working Paper upon the world; if so, Appendix C could deviate some from what I've covered.

Erratica

As I write, diligent reader Dan Saks is in town for the Seattle leg of his North American rock tour. Over a tasty evening repast at the downtown Bellevue Red Robin, we discussed one of my topics from last month: tagged vs. untagged name spaces. Dan brought up a subtlety which I'd failed to mention, but want to touch on here.

Two of the points I raised in the column are that

- C++ implicitly declares tagged type names in both the tagged and untagged type name spaces.
- The untagged name can be hidden by other declarations for the same name at the same scope, leading to the name being used simultaneously as a type and non-type.

In C, to achieve point 1 above, you must use an explicit `typedef`. This has the side-effect of avoiding point 2. Thus, the C `typedef` serves two distinct purposes: it puts the name in the untagged space, and prevents that name from being trumped by another declaration.

Because C++ gives you point 1 for free, you don't need the `typedef`. If, however, you want to avoid point 2, you must use a `typedef` anyway. Moral: In C++, there are still reasons to use `typedefs`, even though that use may appear redundant.

Dan says colloquially that C++ has one and a half name spaces, since the implicit untagged name space entries are so easily masked. Mister Twister, a.k.a. Marc Briand, *CUJ's* Managing Editor, suggests this tagged/untagged business would make great fodder, er, subject matter for a future column. I dunno ... I was instead thinking of a series on abstract declarators. Can't get enough of that sexy topic. □

Notes

- [1] Not to be confused with the Evergreen Point Floating bridge, the proverbial weak link in the Seattle area's freeway chain.
- [2] This function makes gross (but often correct) assumptions about its target character set. For demonstration purposes only. Cars driven by professional drivers on a closed course.
- [3] See my December '95 column for an exegesis on these specifiers.
- [4] (Sung to the tune of the old calypso song): Let us put C and C++ together and see which one is smarter. Some say C, but I say no, C++ got C like a puppet show. Ain't me, it's WG21 that say C is leading C++ astray. But I say it's C++ today smarter than C in every way ...

Four Great C/C++ Tools in One Coordinated Package

- C-Xref for cross-referencing
- C-Tree for function-call analysis
- C-Lines for outlined listings
- C-Format to reformat C/C++ programs

Cross-Reference

```
# - a symbol is declared or defined
& - symbol's address is taken: &abc
= - symbol takes on a new value: i=1; or i++;
! - a function is defined

buf, char * in class String
  s.cpp      8#   14   22=  23   24
                 48   48   48   51   51
dump(void) in class String, returns void
  s.cpp      13#  13!   70   71   73
exit(int), returns void
  c:\include\stdlib.h  81#
  s.cpp      23   44
len, unsigned int in class String
  s.cpp      9#   21=  22   28=
main(void), returns int
```

s.cpp

```
5   class String
6   {
7     public:
8       char *buf;
9       unsigned len;
10    public:
11      void operator += (const String & x
12      String( char *s = 0 );
13      void dump() const
14        { printf( "String = %s\n", buf
15      };
16
17      String::String( char *s )
18      {
19        if( s )
20        {
21          len = strlen(s);
22          buf = new char[len+1];
23          if( !buf ) exit(1);
24          strcpy( buf, s );
25        }
26        else
27        {
28          len = 0;
29          buf = 0;
30        }
31      }
32
33      void String::operator += (const String
34      {
```

C-Vision

Reasons to Choose C-Vision

1. Supports both C and C++ programs
2. Intelligent cross-referencing – can distinguish between definitions, declarations, uses and modifications of variables
3. Intelligent labeling of trees and cross-references provides the types of variables and prototypes of functions
4. Cross-reference and tree information can be dumped to disk as ASCII files and manipulated directly
5. Uses the same syntax engine as the thoroughly tested PC-lint
6. Built with 386 DOS extender technology for high speed and so you won't run out of space.
7. Numerous options for customization of output

Hierarchy Tree

```
1: main(void), returns int, s.cpp @ 64
2: => dump(void) in class String, returns void, s
   |--> operator+(const String &, const String
   |--> print(const char *, ...), returns
   |--> String(const String &), in class St
   |--> operator+=(const String &) in class
   |--> String(char *) in class String,
   |--> operator new[](unsigned int
   |--> exit(int), returns void, st
   |--> strcpy(char *, const char *
   |--> strcat(char *, const char *
   |--> operator delete(void *), re
   |--> strlen(const char *), retu
```

Gimpel Software

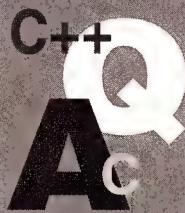
3207 Hogarth Lane, Collegeville, PA 19426

**CALL TODAY (610) 584-4261
Or FAX (610) 584-4266**

Only \$239
Specify MS-DOS or OS/2
30 Day Money-back Guarantee.
PA add 6% sales tax.

C-Vision is a trademark of John Rex and Gimpel Software.
PC-lint is a trademark of Gimpel Software.

Pete Becker



Little-Known Effects of Defining Constructors

Find out why constructors and initializers don't always get along, and when to disobey the gods of OOP.

To ask Pete a question about C or C++, send e-mail to pbecker@wpo.borland.com, use subject line: Questions and Answers, or write to Pete Becker, C/C++ Users Journal, 1601 W. 23rd St., Ste. 200, Lawrence, KS 66046.

QI've been self teaching C++ for about 1 1/2 years now, and I'm comfortable with most of the language features (I mean, it doesn't scare me on a daily basis, as it used to do some months ago). I've tried to compile the following code with two versions of Borland C++ (3.1 and 4.5), and it always give me the error message:

Error test.cpp 19: Objects of type 'C' cannot be initialized with {}

test.cpp:

```
class A {
public:
    A( char c = '\0' ) : ch( c ) {}
    ~A() {}
    // other members...
private:
    char ch;
};

// this is Ok, no complaint
A vowels[5] = {'a','e','i','o','u'};
```

Pete Becker is Senior Development Manager for C++ Quality Assurance at Borland International. He has been involved with C++ as a developer and manager at Borland for the past six years, and is Borland's principal representative to the ANSI/ISO C++ standardization committee.

```
struct B {
    char a, b;
};

// this, obviously, is also Ok
B a_B_struct = { 'a', 'b' };

struct C {
    A a, b;
};

// compiler error !!
C a_C_struct = { 'a', 'b' };
```

The compiler accepts the array initialization for both built-in and user-defined types, calling the proper constructor for the last. With structs, however, I get the error message. The help for the message says: "C++ classes can only be initialized with constructors if the class has constructors, private members, functions, or base classes that are virtual." If so, how can the array be initialized with no error?

I've found no reasonable explanation for this "feature," but it seems to be contrary to the (so-called) "spirit of C++," by presenting inconsistent behavior between the language native types and user-defined types.

Is that a language feature a compiler bug? Am I doing anything wrong? Any hints would be greatly appreciated.

— Marcos Capelini

AYou're right, the compiler doesn't like your struct A in a context where it would be happy with a built-in type. On the surface this compiler behavior is inconsistent. As Emerson reminds us, however, "A foolish consistency is the hobgoblin of little minds..." [1] So let's look

for the reason for this inconsistency, and see if we think it's defensible.

Let's begin by figuring out what the rule here really is. I could start this sort of a discussion off with a citation from the ANSI/ISO working paper, but in reality, I usually do what all of you do: I start off by playing with the compiler to see what it likes and what it doesn't like. Then when I think I've figured out what rules the compiler is applying I go to the authoritative source and try to figure out what rules it says the compiler should be applying. In most cases they're the same.

So, let's start out with something that we know works in C: embedding a built-in type in a struct.

```
Struct s1
{
    int i,j;
};

s1 s = { 1, 2 };
```

As I'm sure you expected, this is perfectly acceptable to my compiler. It's straight C code, initializing a static object.

Now let's replace the built-in type with a struct:

```
struct D
{
    int a;
};

struct s2
{
    D i,j;
};

s2 s = { 1, 2 };
```

This, too, is acceptable, although BC5 gives me two warnings:

```
Warning test.cpp 11: Initialization
is only partially bracketed
Warning test.cpp 11: Initialization
is only partially bracketed
```

If you count lines in my source code, you'll see that both warnings are occurring on the very last line, the one that creates our object of type s2 and initializes it. The warning is legitimate: s2 contains two structs, and the initialization can be

made more specific by adding brackets to delineate the initializers for each of the structs:

```
s2 s = { {1}, {2} };
```

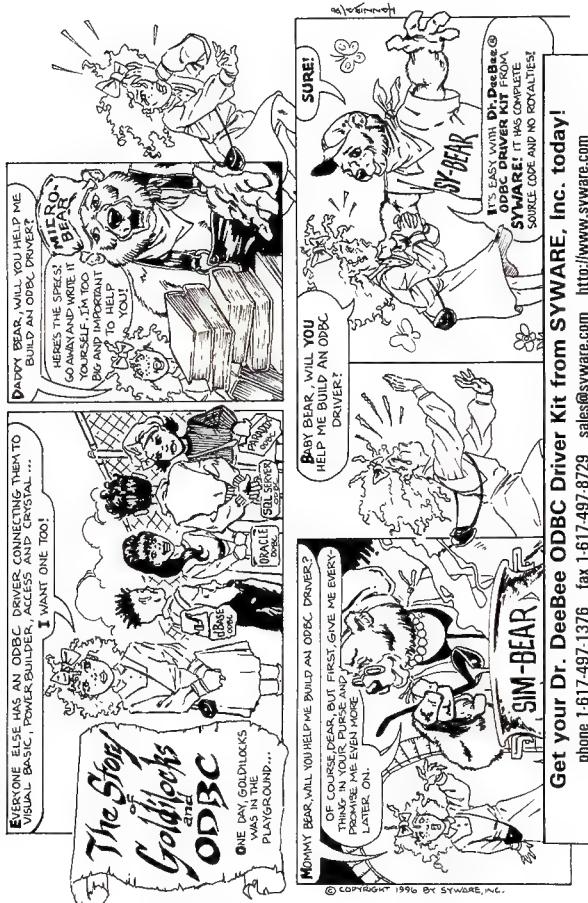
This will quiet the compiler down. The compiler issues this warning because the potential for error is high: when you don't use brackets to set off the initializers for each of the contained structs, you have to count initializers more carefully. Let's change D a bit to see what sort of problems can arise:

```
struct D3
{
    int a;
    int b;
};

struct s3
{
    D3 i,j;
};

s3 s = { 1, 2 };
```

Now the warning is pointing us to a potentially serious problem: the two fields in `s.i` are being initialized, and the two fields in `s.j` are not. That's because in the absence of braces to tell the compiler what to do, the compiler initializes each of the items in the struct



□ Request Reader Service #161 □

in the order it encounters them. The leftover fields are left uninitialized. If we really want to do the initialization this way we can leave it alone. If we want to initialize a member of each of the D3 objects, as we did in the previous examples, we need to use a fully bracketed initialization:

```
s3 s = { {1}, {2} };
```

This initializes `s.i.a` and `s.j.a`, and leaves `s.i.b` and `s.j.b` with their default initialization. Finally, if we want to explicitly initialize all the fields in `s3`, we can do so in two ways:

```
s3 sa = { {1,3}, {2,4} };
s3 sb = { 1, 3, 2, 4 };
```

These two initializers do exactly the same thing. It's a bit harder to tell what's happening with the second form, so I prefer the first. Both are valid in C and C++.

So, the problem your code is having is not because you're using a struct as a member of C. That still works. How is your struct A different from my struct D above? Well, for one thing, it has a constructor. Let's see if adding a constructor to D leads to problems:

```
struct D4
{
    D4(int);
    int a;
};

struct s4
{
    D4 i,j;
};

s4 s = { 1, 2 };
```

Now I get an error:

Error test.cpp 12: Objects of type 's4' cannot be initialized with {}

Does that look familiar? Apparently adding the constructor made the difference. Be careful, though: always dig a little deeper. Let's try a simpler version of this code:

```
struct D4
{
    D4(int);
    int a;
};

struct s4
{
    D4 i,j;
};

s4 s;
```

Now I get a different error:

Error test.cpp 12: Compiler could not generate default constructor for class 's4'

That's right: we've created a class definition that's completely invalid. We got distracted for a moment by our attempts to apply aggregate initialization, and forgot to look at whether we could create one of these things in the first place. So let's help the compiler out by providing a constructor for our outermost struct:

```
struct D5
{
D5(int);
int a;
};

struct s5
{
s5(int);
D5 i,j;
};

s5 s = { 1, 2 };
```

Now we're back to our original problem:

Error test.cpp 13: Objects of type 's5' cannot be initialized with { }

Let's try one more thing: we can get rid of the explicit constructor in s5 if we give D5 a default constructor. Does that work?

```
struct D6
{
D6();
int a;
};

struct s6
{
D6 i,j;
};

s6 s = { 1, 2 };
```

Warning test.cpp 12: Initialization is only partially bracketed

Warning test.cpp 12: Initialization is only partially bracketed

Yes, it works. What if we provide a default constructor that does the same thing as the compiler-generated constructor, that is, it simply calls the default constructor for each of the D subobjects?



Go Virtuoso™ !

DSP and real-time developments on time

DSP

21020

21060

5600x

9600x

320C30

320C31

320C4x

320C5x

Pine DSP core

Oak DSP core

Hyperstone

Other

80x86

3052

68HC11

68HC16

680x0

T4xx

T8xx

80960

ARM

...

Virtuoso's multi-level support

interrupt support

multi-tasking nanokernel

preemptive micrakernel

transparent parallel processing

heterogeneous targets

No compromise on performance

ultrafast (ctxt switch < 1 µs)

small size (0.2 - 6.5 K instr.)

Development support

system generation

(network) loader

debugger and tracing monitor

true portability and scalability

stdio & PC graphics

cross development on PC

MS-Windows, Solaris integration

DSP & multi-media support

drivers

source code

12 months maintenance and

support included

Ease of use

Start on day one !

Virtuoso's multi-tool approach

✓ Real-time O.S.'s :

Virtuoso Micro

Virtuoso Classico

Virtuoso Nano

RTOS implementations :

SP : Single Processor

MP : Multi-Processor

VSP : Virtual Single Processor

✓ Static Schedule generator :

Virtuoso Synchro

✓ DSP Libraries :

Virtuoso Modulo

0, I, II, III, IV, V, VI

✓ Services

Customization and training

Consultancy for HW and SW system design

signal processing

image processing

high speed control

telecommunications

data-acquisition

multi-media

military

space missions

embedded systems

distributed control

process control

Check us out on the web: <http://www.eonics.com>

Eonic Systems, Inc. 12210 Plum Orchard Drive, Silver Spring, MD 20904, USA, Tel. (301) 572 5000, Fax (301) 572 5005, e-mail: info@eonics.com
 Eonic Systems n.v., Lindestraat 9, B-3210 Linden, Belgium, Tel. (+32) 16 62 15 85, Fax (+32) 16 62 15 84, e-mail: info@eonics.com

□ Request Reader Service #162 □

```
struct D7
{
D7();
int a;
};

struct s7
{
s7() : i(),j() {}
D7 i,j;
};

s7 s = { 1, 2 };
```

Error test.cpp 13: Objects of type 's7' cannot be initialized with {}

So it looks like the compiler is happy with aggregate initialization of objects that do not have any user-defined constructors, but won't accept objects that have constructors. We haven't really determined that yet, though: we've only tried default constructors. What happens if we add a constructor that takes an int, so it matches the types that we're using in the initializer list?

```
struct D8
{
D8(int);
int a;
```

};

```
struct s8
{
s8(int);
D8 i,j;
};
```

```
s8 s = { 1, 2 };
```

Error test.cpp 13: Objects of type 's8' cannot be initialized with {}

What about a fully bracketed initialization?

```
s8 s = { {1}, {2} };
```

Error test.cpp 13: Objects of type 's8' cannot be initialized with {}

Let's try the last resort of the desperate: a cast.

```
s8 s = { s8(1), (2) };
```

Error test.cpp 13: Objects of type 's8' cannot be initialized with {}

Still no luck. The rule seems to be ironclad: objects with user-defined constructors cannot be initialized through aggregate initialization. If you look it up in the ARM or in the ANSI/ISO working paper you'll see that that's what it says. You can't do it.

Of course, all this exploration does not get at the answer to the ultimate question: why can't you do it? The logic is quite simple: by adding a constructor to your class you've suggested that the default action that the compiler would have taken would not work correctly. That really shouldn't be much of a surprise: the same thing happens with the default constructor. If you define any constructors in your class the compiler won't generate a default constructor. Add this one to your list of rules: if you define any constructors in your class the compiler won't do aggregate initialization. It believes you when you tell it that it doesn't know enough about your class to do these things right.

[1] "... adored by little statesmen and philosophers and divines." This is from R.W. Emerson, *Essays*, "Self Reliance", (First Series, 1841). I do not, of course, intend to accuse those who seek consistency in programming languages of being in league with "little statesmen."

Q In C++ a class can have public, protected, and private sections. You can give other classes or functions access to the protected and private stuff using friend. Is there an acquaintance? That is, if I said that some other class was an acquaintance of my class, the other class could do public (of course) and protected, but not private.

Actually, I'm pretty darn sure that this doesn't exist now in C++. Do you know if it's going to exist in the future? I know that with some finagling I can accomplish what I want to do. I'd rather do it the easy (but not yet existing) way :-) — Brenda Molina



Bavaria · Germany · Europe · America · Asia · Australia
Phone: ++49-911-20626-0 · Fax ++49-911-20626-18
Stromerstrasse 5 · D-90443 Nuremberg/Germany

□ Request Reader Service #163 □

ANo, there's no acquaintance, and there's not much chance of getting one. There has been no lack of proposed changes to the Standard's specifications of access specifiers, to provide finer control over who can access what. But for the most part, these proposals tend to be workarounds. That is, they attempt to work around design problems in specific programs, rather than provide broadly useful language changes. In such cases, it's better to fix your program design rather than hack the language definition.

The C++ language definition, whether you get it from the ARM or from the ANSI/ISO working paper, does not tell you how to use the things that it defines. That's something you have to figure out for yourself. Used properly, protected members provide a simple, powerful, and controllable mechanism for permitting changes to the behavior of a class. Make sure you understand how this mechanism works, and that you've used it properly, before you decide that it's too limiting.

Public members of a class *C*, of course, can be accessed by any code that can access an object of type *C*. They define the behavior that the class *C* implements. That is part of the contract between the designer and the user of the class *C*. When *C* has a member function named *show*, and *show* is documented to display the object on which it is called, you ought to be able to rely on that behavior in all objects of type *C* and in all objects of types derived from *C*. If you can't rely on that, someone has messed up the class hierarchy.

Protected members are used to help derived classes implement the documented behavior of the base class. Protected members typically do one of three different things: they provide access to important information about the base class that functions in the derived class may need; they provide implementations of the basic operations that derived classes may need; and they provide hooks so that derived classes can modify the details of how the base class behaves.

Let's take the classic example of a base class for objects that can be drawn on the screen:

```
class Drawable
{
public:
    virtual ~Drawable();

    void draw() const;
    void erase() const;
    void move( int x, int y );
protected:
    Drawable( int x, int y );

    virtual void do_erase() const = 0;
    virtual void do_draw() const = 0;

    int get_x_pos() const;
    int get_y_pos() const;
private:
    int x_pos;
    int y_pos;
};
```

C-Index/II
Database Library
Shrink Wrap
Mission Critical
Cross-Platform

Used by leading developers in their award-winning software:
Intuit Symantec Banyan Stratus Boeing
Mindscape Citicorp Timberline FDIC
PG&E Computer Associates Leica JPL

C-Index/II is a C/C++ database library that stores and retrieves application information.

Complete information is available at our web site. To receive the free **C-Index/II** demo disk packet by mail, contact us by phone, fax or email.

For a detailed analysis of how **C-Index/II** will meet your needs, call or write with your specific requirements. A knowledgeable developer will assist you.

Partial Function List

High Level

dbcreate	dadd	dfind
dbopen	dupdate	dseq
mclose	ddelete	dread

Low Level and Multi-User

cdupadd	imgopen	msetlock
cunqadd	imageback	mtstlock
cchange	imagerest	mcrllock
cdelete	imgclose	msetsema
cfind	bcheck	mtstsema
cfirst	bbuild	mcrlsema
clast	addrout	strwrit
cnext	dropout	endwrit
cprev	flush	cispeedon
cnxtrep	cibusnum	cispeedoff

- Windows 3.x, Win95, NT
- Unix Macintosh DOS OS/2
- PharLap Embedded Systems
- Millions of copies in use daily
- Proven reliability since 1983
- All Major Compilers
- Full Source Code Included
- Individual License: \$695

Exclusive Features

- Key Data
- ReadShare
- SpeedRead
- PowerFail Protection

Advanced Features

- Fast B+Tree Indexing
- Unlimited File Size and Key Types
- Complete Portable C Source Code
- Any system with ANSI or K&R C
- Live Cross Platform Data Access
- Multi-User and Single-User Access
- Data and Indexes in Same File
- Multiple Record Formats per File
- Link Directly to App or use DLL
- Over 80 API Functions
- Nine Simple High-Level Functions
- Very Flexible Low-Level Functions



Trio Systems

936 E. Green St. Suite 105
Pasadena, CA 91106

818/584-9706

818/584-0364 Fax

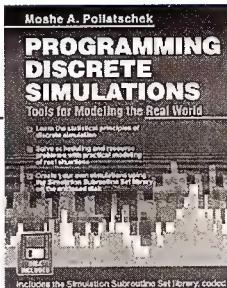
mail@triosystems.com

<http://www.triosystems.com>

A \$10 fee is charged for sending the demo disk packet by airmail outside of North America. C-Index/II, ReadShare, SpeedRead, and PowerFail Protection are trademarks of Trio Systems. All other trademarks are the property of their respective owners.

Unique titles filled with fundamental information and aimed at practical applications.

R & D Books



Programming Discrete Simulations

Tools for Modeling the Real World

By Moshe A. Pollatschek

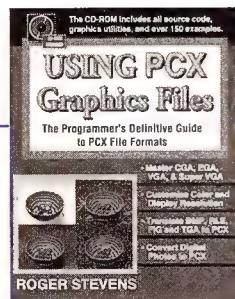
Pollatschek explains clearly how to use discrete simulations to model real world scheduling problems. Whether collecting statistics on telemarketing call volumes, managing factory inventories, or some other scheduling problem *Programming Discrete Simulations* provides the necessary tools. Companion disk includes the Simulations Subroutine Set — a useful set of tools for writing discrete simulation programs.

1995, 350 pp., includes 3-1/2" disk
ISBN 0-87930-449-9

Order #X01 \$39.95

Solve scheduling and resource problems with practical modeling.

Customize PCX graphics to fit your needs.



Using PCX Graphics Files

The Programmer's Definitive Guide to PCX File Formats

By Roger T. Stevens

The most complete how to package available on the widely used PCX graphics file format. Explains techniques for creating and viewing PCX files, also describes in great detail how to manipulate these files. Companion CD-ROM contains all the programs described in the book in both source and executable form.

1996, 450 pp., includes CD-ROM
ISBN 0-87930-432-4

Order #X37 \$49.95

Understanding Self-Similar Fractals

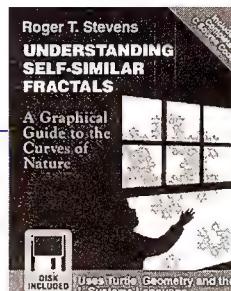
A Graphical Guide to the Curves of Nature

By Roger T. Stevens

Useful to the beginner and expert alike. Stevens provides hands-on lessons in self-similar geometries and algorithms, and a simple to use tool for fine-tuning fractal images. Also explains how to compose entire scenes using a collection of fractal images.

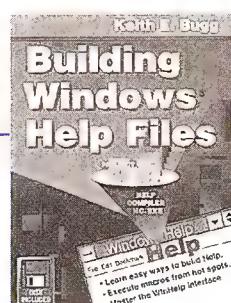
1995, 400 pp., includes 3-1/2" disk
ISBN 0-87930-451-0

Order #V75 \$39.95



Use fractals to depict naturalistic images such as crystals, snowflakes and leaves.

Learn the nuts and bolts of creating Microsoft Windows Help Applications.



Building Windows Help Files

By Keith E. Bugg

A concise programmer's reference for developing the necessary Help files and making them work together. Tells you everything Microsoft forgot to explain about using WinHelp. Companion disk contains code samples.

1995, 280 pp., includes 3-1/2" disk
ISBN 0-87930-439-1

Order #X04 \$29.95



Contact us for a complete catalog or to place your order today!

- call: 1-800-444-4881 or +1-913-841-1631 • fax: +1-913-841-2624
- email: orders@mfi.com • WWW: <http://www.rdboooks.com>

All R&D titles are available in bookstores, or directly from:



An imprint of Miller Freeman

```

void Drawable::draw() const
{
do_draw();
}

void Drawable::erase() const
{
do_erase();
}

void Drawable::move( int x, int y )
{
erase();
x_pos = x;
y_pos = y;
draw();
}

```

The documentation for this class should have two parts, one describing the public interface and one describing the protected interface. The public interface is available to users of the class and classes derived from it. This interface supports drawing a derived object, erasing a derived object, moving a derived object, and destroying a derived object. (Of course, the documentation would describe each of these operations in more detail.) The protected interface is available to designers and implementers of classes derived from Drawable. The protected interface supports creation of an object at a specified position through the constructor. It also enables those with access permission to determine the object's current location, through `get_x_pos` and `get_y_pos`. The protected interface also provides hooks through `do_erase` and `do_draw`, so that the implementer of a derived class can implement erasing and drawing appropriately.

Let's look at a Circle class based on Drawable:

```

class Circle : public Drawable
{
public:
    Circle( int x, int y, int r );
protected:
    int do_draw() const;
    int do_erase() const;
private:
    int radius;
};

Circle::Circle( int x, int y, int r )
    : Drawable(x,y), radius(r)
{

int Circle::do_draw() const
{
draw_circle( get_x_pos(), get_y_pos(), radius, black );
}

int Circle::do_erase() const
{
draw_circle( get_x_pos(), get_y_pos(), radius, white );
}

```

The class Circle adds its constructor to the public interface inherited from Drawable. That is, Circle promises us that we can create a circle at a specified location with a specified radius, that we can draw and erase a circle, and that we can move a circle to a different location. To do these things, Circle uses protected members of Drawable for all three of the basic capabilities that I mentioned earlier. The constructor uses Drawable's constructor to pass the Circle's initial position down to Drawable. The functions `do_draw` and `do_erase` use `get_x_pos` and `get_y_pos` to get the basic data they need from Drawable, and they provide implementations of `draw` and `erase` that are appropriate for a Circle.

Now we can create a Circle and move it around:

```

int main()
{
Drawable *object = new Circle( 0, 0, 10 );
object->show();
object->move( 5, 5, );
object->erase();
delete object;
return 0;
}

```

There's a problem with Circle, though. Suppose you have an object manager that keeps track of Drawable objects, and tries to be smart about not drawing objects that are not currently on the screen. To do that it needs to know the positions of all of the Drawable

QNX® 4.2 Operating System

Realtime 32 Bit POSIX Compliant OS

Multitasking, Multiuser, Peer-to-peer Integrated Networking within a 7K Microkernel

► Development Environment:

**WatCom ANSI C Compiler, 32 Bit
WatCom C++**

► Realtime Scheduler: 3 Options

► Network Support: Arcnet and Ethernet

► Database: Xbase, Btree, SQL and others

► Connectivity: TCP/IP and DOS-QNX links

► Graphics: Direct Graphics and GUI Options

► DOS emulation within QNX supporting Windows 3.1

► Configurable for Embedded Applications

©QNX is a registered Trademark of QNX Software Systems Ltd.

Call or write for more information:



COMPUTER PRODUCTS, INC

P.O. Box 14010 Tulsa, OK 74159 USA
Tel: 918/742-1816 Fax: 918/749-7113
VISA/MC

□ Request Reader Service #166 □

objects. That seems simple enough: just have the object manager call `get_x_pos` and `get_y_pos` whenever it needs to know where an object is. The problem is, it can't, because these functions are protected. Your first reaction might be to make the manager class a friend of `Drawable` so that it can call these functions. But this is overkill, since manager will then have access to `Drawable`'s private data as well. Wouldn't it be nice if you could say that the object manager can get at `Drawable`'s protected members but not its private members? Alternatively, wouldn't it be nice to let the object manager call `get_x_pos` and `get_y_pos` but no other protected or private members?

Tinkering with access rights might seem like a good answer at first glance, but there's a much better solution: make `get_x_pos` and `get_y_pos` public. After all, you will almost certainly encounter other situations in which you need to find the location of a `Drawable` object. Make this change and the problem goes away, and we have a class with a broader and more usable public interface.

No doubt some of you will object at this point that I've set you up: making `get_x_pos` and `get_y_pos` protected instead of public is obviously wrong; you suspected it all along. I plead guilty, but with an explanation: it's easy to get so focused on the details of making your code work that you lose sight of the overall design. When you run into an access problem, think about your class's public interface and protected interface. Don't just add friend declarations or wish for fancier options for controlling access. Look at how you got into your current situation, and consider all the possible ways out. Often

the best solution is to rethink the interfaces to your class. If the design is wrong, fix it.

Q We are trying to develop a Windows application that has multi-lingual support. To do this we have implemented a `Msgs` class in C++ that reads in a text file as an array of strings (`TArray<string>`). Since almost every subclass in the application needs to generate error/information messages it seems inefficient to pass the `Msgs` instance down the tree to make its members accessible. Declaring the `Msgs` object as a global would break OO rules (wouldn't it?).

A colleague has tried using a pointer to the parent class containing the `Msgs` object. Since this pointer has a type defined by the parent then the class using the `Msgs` object needs to know that type, thus again breaking the OO rules.

Inheritance could be used but the subclasses are not forms of `Msgs`, but have `Msgs` as part of them. For example:

```
class Msgs
{
private:
    TArray<string> tasMsgs;
public:
    string& operator[](int index);
};

class OtherClass
{
public:
    void error_out(int id)
    {
        cout << tasMsgs[id]; // an error since
                            // OtherClass doesn't
                            // know about tasMsgs
    }
};

class Parent
{
public:
    Msgs mTheMessages;
    OtherClass oAnObject;
};
```

Any suggestions? — Tom Robinson

A There's a saying among those who trek the back country: if the map doesn't match the terrain, trust the terrain. The same thing applies to the "rules" that we use to suggest good design and coding techniques. If the rules don't fit your application, make sure the application works right.

You've mentioned two "rules" of object-oriented programming: global variables are bad and class interdependencies should be minimized. Both are generally good rules; both should also be violated when appropriate, and not taken as absolutes.

Sometimes a program uses data that is truly global. If that's the case, make it global. That's often the case with error messages. Sure, you could define each message in the place where it's needed, but then you run into exactly the problem you're trying to solve:

Free from Microsoft!

The Developer's Black Book

Full of tips and useful contact information, this hot resource was created by Microsoft with you in mind. Designed specifically for professional developers of Windows-based applications, this handy reference is available now on newsstands with the September issue of *Microsoft Systems Journal*.

To boost your programming prowess in Windows add *MSJ* to your development toolbox.

Get your copy today. On Newsstands Now!!!

when you want to localize the application (that is, translate messages into a language appropriate for your customers), it's much easier to have all the messages together in one place. Let's take a look at your code rewritten with the messages stored globally:

```
class Msgs
{
public:
    Msgs();
    string operator[](int index) const;
};

const Msgs ErrorMessages;

class OtherClass
{
public:
    void error_out(int id)
    {
        cout << ErrorMessages[id];
    }
};
```

I've made several changes to the sketch of your class `Msgs`. I added a declaration of a default constructor; I assume this was part of your original class, and you left it out of your note to make the class definition shorter. I added the default constructor back in because it's important to remember that it's there: it's responsible for loading the messages from wherever they're stored, presumably in some file that the program knows about. I also removed the actual array of messages. When we're talking about program design we're not interested in the internal details of how a class does whatever it does. All that matters is that you can index into an object of type `Msgs` and get back a string. I changed the return type of the index operator from `string&` to `string`, and I've added a `const` qualifier to it, to indicate that using it won't modify the object that you're calling it on. That means that the object `ErrorMessages` which I've added can be a `const` object.

Maybe you've noticed my change puts a heavy emphasis on `const`-ness. There's a reason: one of the biggest dangers in using global data is getting confused about which functions change the data. You may find that you no longer know whether the object actually contains what you think it contains. If you forbid your code to modify your global data you can have a great deal more confidence in its integrity. Once the `ErrorMessages` object has been initialized no one is going to change it. Your messages are safe.

Even when making things like `ErrorMessages const`, you can still carry out important optimizations such as lazy evaluation. Suppose you don't want to take the time to read the message file at startup, but instead want to wait until the program needs to display an error message. You can do that and still insist that `ErrorMessages` is a `const` object. We refer to objects like this as *logically const*: calling any member functions on `ErrorMessages` produces the same results, regardless of the preceding sequence of calls to its members. You'll always get the same message when you look up the same index.

To implement this behavior, keep a flag inside `Msgs` to indicate whether you've read the message file, and check that flag in the

index operator. Both the flag and the message array will be modified by the index operator on its first pass. Since I've marked the index operator as a `const` operation, we have to tell the compiler that it's okay for `const` member functions to modify these two data members. That's the job of the keyword `mutable`.

```
class Msgs
{
public:
    Msgs() : data_present(false) {}
    string operator[](int index) const;
private:
    mutable bool data_present;
    mutable Tarray<string> messages;
};

string Msgs::operator[](int index) const
{
    if( data_present == false )
    {
        // insert code to read message file
        data_present = true;
    }
    return messages[index];
}
```



CDROMs!

Call today for your free catalog of all our CDROM titles!

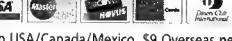
February 96 Version! C User Group CDROM

The new version of this popular CDROM gives you the entire C Users Group Library, a collection of more than four hundred MB of user supported C source code. You get function libraries, text editors, compilers, text formatters, interpreters, BBS's, etc. \$ 49.95

	Official Slackware™ Linux – Slackware 3.0 Linux (ELF), by author of Slackware. Learn Unix! Easy install. X Window XFree86 3.1.2. Total dev. environment: Tcl, GCC (2.7.0) compiler. Doom & Abuse games. 2 discs \$ 39.95
	Slackware Subscription – convenient about every 3 month update, each \$ 24.95
	FreeBSD® 2.1 – Berkeley BSD for PC. Rock solid Unix-like, easy install, full src, XFree86 3.1.1, dev. tools, etc. (Every 6 mo. subscription \$24.95) \$ 39.95
	POV-Ray® 3-D Raytracing : 1000 wow! images, scene src & binaries. \$ 39.95
	Avalon – 3D objects, graphics source, libraries, graphics documents. \$ 39.95*
	CICA® MS Windows – very popular. 5000 current shareware. French / Italian German / Spanish translations. 2 discs. (3 mo. sub. \$19.95 issue.) \$ 29.95*
	Hobbes OS/2 Archived – OS/2 Magazine's 1994 Ed's choice . 600 MB selected OS/2 applications, drivers, docs. (3 mo. sub. \$19.95) \$ 29.95*
	Simtel® MSDOS – Classic MSDOS programming, tech shareware. 2 disc. \$ 29.95*
	Science Library – Tech, engineering, science shareware + index book. \$ 39.95*
	GNU – source & compiled for FreeBSD, Sun OS 4.1.3, Solaris 2.x. (1/96) \$ 39.95
	Perl – Source and binaries for many systems, documents & sample code. \$ 39.95
	TEX – TeX and LaTeX. 1000 MB of code, binaries, utilities & docs. \$ 39.95
	Ada – Ada compilers, tools, source code, and documents. Two discs. \$ 39.95
	Math Solutions – Math applications, function libraries, source code, docs. \$ 39.95*
	Clipper – Database language source, libraries, utilities & docs. \$ 39.95*
	Music Workshop – Music demos, players, wav. (MS Windows) \$ 39.95*
	Internet Info – 15000 computer and internet docs. IENs, RFCs, FAQs. \$ 39.95*

*Shareware programs require separate payment to authors if found useful
— Authors of similar discs wanted. Please email discdev@cdrom.com —

ORDER NOW! 1-800-786-9907

Shipping is \$5 in USA/Canada/Mexico, \$9 Overseas per order

Walnut Creek CDROM
4041 Pike Lane, Suite D-235 Concord, CA 94520
+1-510-674-0783 • FAX: +1-510-674-0821 • email: orders@cdrom.com • <http://www.cdrom.com/>

All of our CDROMs are **unconditionally guaranteed!**
adcode: 235

The other risk in using global data is that two or more modules may use the same name for different global objects. That's easily solved: just put the global objects into a *namespace*:

```
Namespace Errors
{
class Msgs
{
public:
    Msgs();
    string operator[](int index) const;
};

const Msgs ErrorMessages;
```

The construct

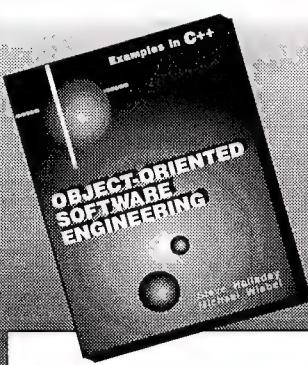
```
Namespace Errors
{
    ...
};
```

defines everything between the curly braces to be within the Errors namespace scope. To access names defined in this scope, code outside the scope must add the namespace name as a qualifier:

```
class OtherClass
{
public:
    void error_out(int id)
    {
        cout << Errors::ErrorMessages[id]; //qualifier req'd
    }
};
```

So far I've been assuming you have only one set of error messages to be used by all classes in your application. The techniques I've described work fine in that situation. If you want to break your error messages down into smaller groups to make them more manageable or more flexible you've got do to a bit more coding. Let's say you've designed your class OtherClass to be used as a member of two separate classes, Class1 and Class2. You want it to display one set of error messages when it is used in Class1 and a different set when it is used in Class2. This situation is fairly common, for example when OtherClass is some sort of helper class, providing services to Class1 and Class2. In that case, the program must report a failure in the helper class in terms that make sense for Class1 and Class2 respectively. You must be able to tell OtherClass which set of messages to use, and that means that you must give it a pointer or reference to a Msgs object holding the appropriate messages. Here's a sketch of how this would be done:

IMPROVE YOUR SOFTWARE ENGINEERING WITH OBJECT-ORIENTED METHODS



OBJECT-ORIENTED SOFTWARE ENGINEERING
by
Steve Halladay & Michael Wiebel

Learn how to:

- Engineer applications to minimize costs
- Manage all six phases of the software lifecycle
- Use recursion in the object-creation process

USE RECURSION FOR PRECISE DESIGN MANAGEMENT

This book demonstrates object-oriented principles for each phase of development, from specification through maintenance. A comprehensive example coded in C++ ties the demonstration together. For precision, the design method is presented in a pseudocode recursive algorithm.

ORDER TODAY!

Order: Book T37 Object-Oriented Software Engineering

913-841-1631
FAX 913-841-2624



Enhance and Analyze Images on ANY PC!

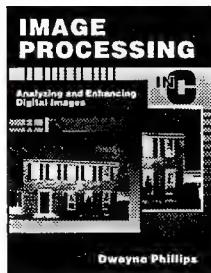


Image Processing in C
by Dwayne Phillips

Learn how to manipulate, analyze, and process images with CIPS (C Image Processing System) on any PC from a 286 to a Pentium. Phillips has distilled image processing to its essentials and clearly explains image analysis and enhancement techniques. You can use CIPS to read and write TIFF images; display them on monochrome, CGA, VGA, and EGA; and print them on a laser printer, graphics printer, or character printer.

Manipulate and Process Images with CIPS

The long-running *C Users Journal* series is now a book with updates for the TIFF 6.0 specification and entirely new chapters on Boolean operations and a batch file manager. Now you can manipulate images without needing a special image processing board, math coprocessor, or top of the line display!

ORDER TODAY!

Use code T60C when ordering
Image Processing in C

913-841-1631
FAX 913-841-2624



```

class Msgs
{
public:
    Msgs(const char *);
    string operator[](int index) const;
};

class OtherClass
{
public:
    OtherClass( const Msgs& msg)
        : ErrorMessages(msg) {};
    void error_out(int id)
    {
        cout << ErrorMessages[id];
    }
private:
    const Msgs& ErrorMessages;
};

class Class1
{
public:
    Class1() : ErrorMessages(
        "Class1msg.txt" );
    other(ErrorMessages) {}
private:
    const Msgs ErrorMessages;
    OtherClass other;
};

class Class2
{
public:
    Class2() : ErrorMessages(
        "Class2msg.txt" );
    other(ErrorMessages) {}
private:
    const Msgs ErrorMessages;
    OtherClass other;
};

```

In every instance of `OtherClass` we've replaced the global variable `ErrorMessages` with a reference to the object that holds its error messages. Now, whenever we create an object of type `OtherClass` we specify the object that it uses to get its error messages. This technique gives us the most flexibility in structuring our error messages. Note that we've done this without requiring `OtherClass` to know anything about the structure of `Class1` or `Class2`. Rather, those classes know their own structure, and know that they need to give an object of type `Msgs` to `OtherClass` for its mes-

sage handling. `OtherClass` has to know only about `Msgs`, which is information that it already had anyway.

You're right to be concerned about global data and about structural interdependencies. Don't overdo it, though. One of the biggest challenges in program design is to recognize when it's all right to break the rules. □

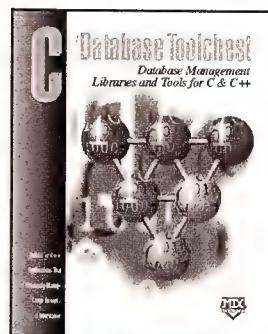
Organize Your Data for Quick and Easy Access

For programs that manage large amounts of data, you need tools that can efficiently organize and access that data. The **C/Database Toolchest™** provides these tools for your C and C++ programming projects.

Both C and C++ libraries are included so you can choose the most appropriate interface for your applications. Whether you choose C or C++, the interface is very simple to use.

The libraries provide a complete set of database management functions. More than 150 functions support file operations such as create, open, and close; and record operations such as add, find, modify, and delete. Example programs clearly illustrate how each function is used.

In multi-user or network environments, multiple programs can share the same database. An entire database, or individual records may be locked to temporarily block access by other programs.



Records may contain any number of string, numeric, and binary fields. Numeric values may be stored in string or binary format to maximize portability or efficiency. Strings are automatically stored as variable-length fields to minimize the amount of disk space required.

Indexes give you control over how your data is organized and accessed. Fast record access is achieved through an advanced B+-tree indexing algorithm. Any number of indexes may be created from any combination of record fields.

The LDM database manager provides an easy way to interactively create and edit databases. Also included are tools for viewing and compressing databases, and for converting to and from dBASE® format data files.

Many combinations of compilers and operating systems are supported. Since all of the C and C++ source code is included, you can build new libraries with just about any standard C or C++ compiler.

**60 DAY
MONEY-BACK
GUARANTEE**

To Order Please Call:
1-800-333-0330

For Technical Questions:
Tel: 1-214-783-6001
Fax: 1-214-783-1404

 Mix Software

1132 Commerce Drive
Richardson, TX 75081

Only \$49.95

Includes All C and C++ Source Code

Select Operating System Version(s):

C/Database Toolchest for DOS/Win3.1 Win95/NT OS/2 \$49.95

Add \$10.00 for each additional version (2 versions for \$59.95, all 3 for \$69.95)

Name _____ Shipping (\$5 US, \$10 Canada, \$20 Foreign) _____

Company _____ Subtotal

Street _____ Texas Sales Tax (.0825 x Subtotal)

City _____ Total (Subtotal + Texas Sales Tax)

State _____ Zip _____ Cash, Check or Money Order Enclosed

Country _____ Paying by Visa, MC, Amex, Discover Card

Telephone _____ Card # _____ Exp. _____

Software supplied on 3.5" disks (1.44Meg HD Format) • Additional operating systems may be supported.

□ Request Reader Service #168 □

Write for CUJ!

Do you sometimes get the feeling your talents go unnoticed? Do you feel like a cog in the software machine? Put some spark in your career, by writing an article for CUJ. Writing for CUJ is rewarding, pays good money, and gets you recognition for your hard work.

Fire Up the Word Processor

Don't wait for an English degree. You needn't be a professional writer to write for *CUJ*. *C/C++ Users Journal* is written by practicing programmers, *for* practicing programmers. If writing is not your strong suit, we have a friendly and experienced editorial staff to help you over the rough spots. Our only requirement is that you have something worthwhile to say. If you've developed a useful technique or tip you want to share with other programmers, we want to hear from you, and so do our other readers.

Write for Real Programmers

CUJ is written for real programmers just like you, who want to see practical treatments of real programming problems, or tutorials that make theory or complex issues more accessible. Your article should concern C or C++, and should provide examples of working code wherever possible. We also publish Book Reviews, User Reports, and very short articles (< 1000 words).

Solid Beats Pretty

When it comes to code, we're picky about the things that count. We're not fussy about where you put your curly braces; we do care about good design and well-crafted code. While we recognize that occasional bugs are inevitable, we expect would-be authors to make a diligent effort to eliminate them before submitting an article.

How to Submit an Article

If you have an idea for an article, send us a proposal: a brief abstract and outline, and a brief resume of your qualifications. Ask for a copy of our Author Guidelines for guidance in writing a manuscript. You must submit your manuscript in electronically readable form. Files must be in raw ASCII format, flush left, with one empty line between paragraphs. Submit proposals, manuscripts, and Author Guideline requests to:

Marc Briand
C/C++ Users Journal
1601 W. 23rd St., Suite 200
Lawrence, KS 66046-2700

PH: (913)-841-1631
FAX:(913)-841-2624
e-mail: marc@rdpub.com
mbriand@mfi.com

Each month we print both theme and non-theme articles. Any time is a good time to submit an article to *CUJ*, but to print it as a theme article we must receive it by the deadlines shown on our Editorial Calendar. Check out our Editorial Calendar on this page if you're looking for article ideas.

1996 & 1997 Editorial Calendar

Issue	Theme	Drafts Due
December	Optimization and Efficiency	16 Aug 1996
January '97	Cross-Platform Development	13 Sept 1996
February '97	Real-Time/Embedded Systems	11 Oct 1996
March '97	UNIX	8 Nov 1996
April '97	Database/Persistent Objects	6 Dec 1996
May '97	Internet Development	3 Jan 1997

Upcoming Themes

Optimization and Efficiency. However powerful computers become, efficiency is always an important issue. Serious programmers need tools for determining where a program spends its time, and reports that highlight how a program consumes storage. They also need techniques for making programs faster and smaller, with a minimum loss of readability and maintainability. We are looking for articles on how better to measure C and C++ programs, and how to improve their performance. We strongly prefer quantitative measures to qualitative claims of improvement.

Cross-Platform Development. We are looking for articles that solve tough problems in cross-platform development, and deal with topics such as the following: C++ classes that encapsulate system-dependent interfaces; improving portability of floating-point code; small libraries that provide useful services across many implementations; tools that check for nonportable constructs; judicious use of Java with C/C++ for platform independence; comparing STL bugs and capabilities across platforms; creating cross-platform persistent objects; platform capability test suites.

Overworked topics (require fresh perspective): Little vs. Big Endian problems.

Real-Time/Embedded Systems. Possible topics include: how to adapt (pervert) popular hosted C/C++ compilers to produce embedded code; testing embedded software on conventional systems; networking embedded systems; using Unix/DOS/Windows for data acquisition and control; C++ on embedded systems; performance prediction, estimation, and/or measurement on real-time systems; failsafe mechanisms for mission-critical code; controlling peripheral devices from C/C++ programs.

Special Request

Templates are standard in C++, but today's compilers all implement them differently. We'd like to see a comparison of template handling abilities in the current crop of popular compilers — differences in syntax, instantiation, what you can and can't do, and implications of the above.

Other Topics of Interest

Java	Multithreading Algorithms
Memory Management	Tips 'N' Tricks
X-Windows Exceptions	OOP

Victor R. Volkman



STL Help on the Web

**You don't have to be a C++ expert to use STL.
The World Wide Web is a great place to find
tools, tutorials, and moral support.**

Focus on STL

The WWW has historically been one of the most informative sources of information about leading-edge programming technologies. Though STL has been around for a few years now, it has only recently made its way out of computer science labs and into the toolkits of corporate application developers. My quick survey of WWW via the AltaVista search engine yielded more than 1,000 documents. After considerable weeding, I'm happy to share with you some of the more relevant pages in several sub-categories: tutorials, vendor-specific pages, and pointer-pages:

STL Tutorials

Jak Kirman <jak@cs.brown.edu> of the CS faculty at Brown University has thrown together what he considers "a very modest tutorial" on STL. You can read it as a hypertext document or print it out in Postscript for offline perusal. Kirman motivates the discussion with a look at the STL sort template while introducing each of the five categories of STL components one at a time. Other sections address the philosophy of STL and suggest how to extend STL to meet your own needs. Kirman makes good use of graphical figures to illustrate his points. Check it out for yourself at:

<http://www.cs.brown.edu/people/jak/programming/stl-tutorial/tutorial.html>

Victor R. Volkman received a BS in Computer Science from Michigan Technological University. He has been a frequent contributor to the *C/C++ Users Journal* since 1987. He is the author of the book *Windows Programming with Shareware Tools*. He can be reached at the HAL 9000 BBS (313)663-4173, URL <http://www.HAL9K.com/home.htm>, or email to sysop@hal9k.com.

Mumit Khan <khan@xraylith.wisc.edu> at the Center for X-Ray Lithography at the University of Wisconsin provides an exhaustive amount of advice and coding examples (163Kb) in his "STL Newbie Guide." Major headings include: Writing Container Objects, Pointers and STL, STL Iterators, Persistent STL, and a new section about Predicates, Comparators, and General Functions. In addition to how-to information, he specifically addresses common pitfalls, such as "Gotchas with Storing Pointers to Objects." Its all at:

http://www.xraylith.wisc.edu/~khan/software/stl/STL.newbie.html#class_howto

David R. Musser <musser@cs.rpi.edu> has collected some introductory information, example applications that use STL, and pointers to other STL reference sources. Specifically, Musser's "Standard Template Library" page includes example source code for demonstrating STL vectors and vector iterators, an anagram checking application, and searching and sorting user-defined records. More advanced users may want to look at the official ANSI/ISO STL specifications (available in Postscript form) there.

<http://www.cs.rpi.edu/~musser/stl.html>

David R. Musser and Atul Saini, authors of *STL Tutorial and Reference Guide: C++ Programming with the Standard Template Library* [see a review of this book elsewhere in this issue — mb] offer a set of source code resources specific to their book. More than 70 example programs are available and you

can browse selected parts of the book. Certainly, the website will help you decide whether Musser and Saini's book would be appropriate for you.

<http://www.aw.com/cp/musser-saini-source.html>

STL Vendor Specifics

G. Bowden Wise <wiseb@cs.rpi.edu>, of the CS faculty at Rensselaer Polytechnic Institute, has published a pragmatic set of notes entitled "Using the The Standard Template Library with Borland C++." In this document, he describes how to make STL work properly with Borland C++ 4.5 and optionally with Borland's Object Windows Library (OWL). Specifically, he details required changes to algobase.h, function.h, defalloc.h and other header files. Also appearing are detailed explanations of the known warnings that the Borland compiler emits while processing STL. You can read about it at:

<http://www.cs.rpi.edu/~wiseb/stl-borland.html>

or download Wise's patched STL files directly at:

<http://www.cs.rpi.edu/~wiseb/stl/stlexp.zip>

Microsoft's "MFC and Standard Template Library (STL) Alert" page provides a blistering disclaimer in which it places sole responsibility for STL back on Hewlett-Packard (who originally made STL freely available, but retained the copyright). Furthermore, it goes on to say that Microsoft Product Support Services will not assist in technical support issues on STL. Nevertheless, it does highlight required changes to header files (involving namespaces) that you can make yourself to achieve compatibility with the Windows SDK or Microsoft Foundation Classes (MFC) and STL. Visit them at:

<http://www.microsoft.com/visualc/v4/v4tech/stlchg.htm>

Warren Young <tangent@cyberport.com> publishes the "STL Resource List," an extensive online catalog of pointers to pages detailing the STL support of C++ compilers. It also contains pointers to FTP sites where you can download a specially-tweaked version of the STL files for your specific compiler. Among the STL vendors cited are:

- Symantec C++ 7.x for Windows and 8.x for Macintosh
- Metrowerks C++ 7.0 (for Macintosh)
- Watcom C++ 10.0
- GNU C++ 2.7.1 and 2.7.2
- Rogue Wave's tools.h++ (enhanced version of STL)
- ObjectSpace's STL <Toolkit> (a multi-platform STL)

 <http://www.cyberport.com/~tangent/programming/stlres.html>

STL Pointer Pages

Here's some "pointers to pointers." These pages provide a rich set of jumping-off points for your own exploration of STL:

 <http://www.cs.rpi.edu/~wiseb/stl-notes.html>

<http://www.xraylith.wisc.edu/~khan/software/stl/STL.newbie.html#resources>

<http://cs.ucr.edu/~jcharvat/stl.html>

STL Usenet Newsgroup Resources

Usenet newsgroups provide an ongoing discussion about the pros and cons of STL implementations. Jan Charvat, a graduate student in Computer Science at the University of California (Riverside), recommends the following newsgroups for STL coverage:

- comp.std.c++
- comp.lang.c++
- comp.lang.c++.moderated

Before posting to Usenet, make sure you've read the STL FAQ, available at:

 <ftp://butler.hpl.hp.com/stl/stl.faq>

Reader Suggested Links

Paulo Eduardo Neves <neves@lmf-di.puc-rio.br> at Laboratorio de Metodos Formais in Rio de Janeiro, Brazil submits his favorite set of C/C++ links. As usual, I'll provide a brief description of each.

The Code Economics project (CodEc) provides a useful set of links to freeware linear algebra (matrices) and numerical analysis tools for C/C++. It also hosts the source and documentation for Chris Burchenall's MatClass and Robert Davies NewMat08. The MatClass C++ matrix class library includes makefiles for Borland and Microsoft C++. The NewMat08 matrix class library supports

**Build a better Windows app
with WDJ code and hands-on
technical information!**



**Articles — over 300 in all —
from every issue of *Windows
Developer's Journal***

Complete, reusable source code

Powerful full-text search

Flexible, information-rich interface

Imagine — four years worth of the **best** Windows programming information available anywhere and the search tool to find exactly what you need! That's what you get with the *Windows Developer's Journal CD-ROM*.

The *WDJ CD-ROM* puts the technical information from *Windows Developer's Journal* on your desktop in a useful, timesaving format — no more fumbling through piles of old magazines to find that key article on VxDs or NT security. You can find any article, on any topic — and export the relevant code to your computer, instantly. Thinking about buying a Windows book? Browse Ron Burk's succinct and honest reviews to find the good ones. Compiler bug bite you? Check out Bug++ of the Month. Into NT? Pull up Paula Tomlinson's "Understanding NT."

Win95, NT, MFC, VxDs, memory management, debugging — if it involves Windows programming, you'll find it on the *Windows Developer's Journal CD-ROM*.

\$49.95
(plus shipping
and handling)

**Miller Freeman, Inc.
1601 West 23rd St., Ste. 200
Lawrence, KS 66046 USA
www.wdj.com**

Windows DEVELOPER'S JOURNAL CD-ROM
DECEMBER 1991 - DECEMBER 1995

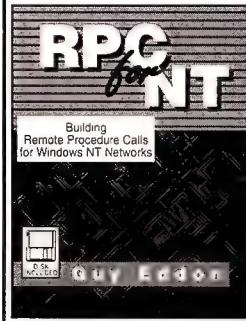
**Phone: 800-444-4881
913-841-1631**

Fax: 913-841-2624

Email: orders@mfi.com

**Order Yours TODAY!
1-800-444-4881**

**See What an NT Server
Can Really do!**



RPC for NT
Building Remote Procedure Calls for Windows NT Networks

Build Remote Procedure Calls from basic to Sophisticated

Guy Eddon guides you from a simple "Hello World" RPC application through a series of increasingly demanding programs that encompass all aspects of remote procedure calls. You'll learn about:

- Implicit and explicit binding procedures
- Structured exception handling
- Multithreaded servers
- The RPC name service

Order your copy today!

Use code T58C when ordering *RPC for NT* with disk included

\$39.95
plus shipping

R&D Technical Books

913-841-1631
FAX 913-841-2624

Borland, Microsoft, AT&T, Zortech, Watcom, GNU, and most other C++ platforms.

 http://netec.mcc.ac.uk/~adnetec/CodEc/C_Cpp.html

Custom Innovative Solutions Corporation publishes its C source code in shareware form on the web. You're free to look at and download the source code. If you find the code useful in a commercial software setting, you are asked to pay a small licensing fee. Specific categories include: CGI Tools, File Handling Utilities, Graphics, Indexing and Searching, Numeric Conversions, Parameters and Configuration, Patterns and Parsing. I was most impressed by CNVL2.C, which can convert a long integer to a printable representation in any numbering base (from base 2 to base 90).

 <http://www.cisc.com/src/c/main.html>

The online hypertext version of the GNU C Library reference manual is very convenient to use. You have instant access at your fingertips using the three-tiered table of contents to a variety of topics. Here's just a sampling of the topics:

- Pipes and FIFOs
- Searching and sorting
- Pattern matching and shell-style word expansion
- Calendars, timezones, and alarms
- Defining and using signal handlers
- Locales, internationalization, and extended character sets

 http://www.ia.pw.edu.pl/tex-info/libc/libc_toc.html#SEC492

Ajay Shah <ajayshah@cmie.ernet.in> at the Centre for Monitoring Indian Economy, Bombay, compiles the "Free C/C++ Sources for Numerical Computation" page. As the name implies, it is an index of pointers to other types of information, including:

- Free source code available on the net
- Books which come with source code, and hence act as low-cost libraries
- Articles and documents, especially those available over the net

Unlike many other resource lists, this one specifically cites the author name, supported platforms, FTP site, and version information.

 <http://gauge.phys.uva.nl:2001/c-sources.html>

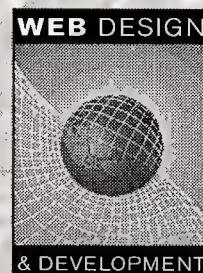
or try

 <ftp://ftp.usc.edu:/pub/C-numanal/numcomp-free-c.gz>

The "C++ Virtual Library," maintained by Lutz Lilje <lilje@desy.de> is a website that has something for everyone from the C++ newbie to those on the bleeding edge of the C++ standards. As you might expect, such a site is primarily focused on providing useful links to authoritative sources of information. Lilje's

Friends don't let friends build static web pages.

<http://www.web96.com>



Web Design & Development '96 East
October 28–November 1, 1996, Washington, D.C.

page is one of the most active programming resources on the web and has earned the coveted "Four-Star Site" award from Magellan for his efforts. Specific areas of interest include:

- Getting Started: Documents and sources on C++ and OOP
- Learning C++: Virtual courses and tutorials
- Free Packages: Freely available C++ packages of all kinds
- Conferences: List of OOP and Computing conferences
- OOLP: Discussion on Object-Oriented Literate Programming
- Tools & Products: Descriptions of products
- General OO: Object-Oriented programming resources.

There are few sites I visit often enough to merit a bookmark, but the "C++ Virtual Library" is definitely on my hot links list.

 <http://info.desy.de/user/projects/C++.html>

If you would like to share some of your favorite C/C++ sources, please email me at sysop@HAL9K.com.

C/C++ Mailbag

Here's a sampling of queries which have come to me recently via e-mail.

Larry_Siden@msn.com writes:

At my work, we might be able to use a program that can analyze C source code and generate a call tree along with various programming metrics. Do you know of any shareware like that? There's commercial stuff, but it costs in the thousands. I'd like to find something more reasonably priced.

C/C++ Answer Man:

There's no need to pay a fortune for good source code analysis (how about just \$47 for a great shareware program?). Try C Exploration tools for Windows. Its C Function Tree and C Structure Tree make program and data structure logic as clear as can be. As for metrics, I'll highlight only a few of them here:

- Number of source code lines for every file
- Number of included files for every source file
- For every defined function:
 - Number of lines
 - Number of functions called
 - Number of flow control statements
 - Maximum brace nesting level
 - Whether it's used only in this file
- Function or data type reference list for every file
- Location of multiply defined functions and data types
- Location of all overloaded C++ functions
- #include file dependencies for every source file
- #include file tree for nested #includes
- Cross-reference for every function or data type
- Parent/child relationship for functions & data types
- C++ class inheritance tree
- C structure/union byte offset calculation

Check out volume #437 on the CUG CD-ROM or by FTP at

<ftp://ftp.cdrom.com/pub/cica/win3/programr/cxtw107.zip>

75300.2257@compuserve.com writes:

I would prefer a better method of searching for information and code on the CUG CD-ROM (please bear in mind that my only exposure is the 1994 CD). A hyper-text style index with search functions is preferable.

C/C++ Answer Man:

Ask no more! The June 1996 edition of C/C++ Users Group CD-ROM includes a hypertext index of volumes 400-437, by Subject keywords, Title, Language, and Platform (OS/CPU). Whats more, you even get a free licensed version of the popular I-VIEW offline HTML viewer (<http://www.talentcom.com/iview/iview.htm>) to use with it. A hypertext index of CUG CD-ROM is available online at <http://www.HAL9K.com/cug/>.

d8342012@ccunix.ccu.edu.tw writes:

I am interested in simulation, especially discrete event-driven simulations.

C/C++ Answer Man:

Well, since you didn't specifically ask for C or C++, I'll assume that C++ class libraries will work for you. I have two specific packages I think you should look at: C++SIM and CNCL.

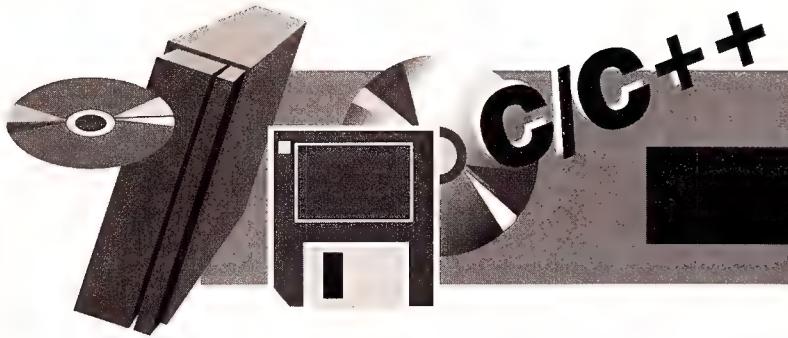
The C++SIM discrete event process based simulation package provides Simula-style class libraries. C++SIM was written by M.C. Little and D. McCue at the Department of Computing Science in the University of Newcastle upon Tyne (England). The same distribution also includes the SIMSET linked list manipulation facilities. According to MacLennan (1983), Simula was the first computer language to incorporate the ideas of "class" and "object" constructs in 1967. C++SIM currently claims usability only on UNIX workstations, such as SUN Sparcs. C++SIM version 1.0 (released 06/15/92) is now available as volume #394 on the CUG CD-ROM. See also the C++SIM home page:

<http://ulgham.ncl.ac.uk/C++SIM/homepage.html>

The Communication Networks Class Library (CNCL) is a C++ library created by the Communication Networks department of Aachen, University of Technology, Germany. CNCL is both a class library featuring generic C++ classes as well as a simulation library with strong points in random number generation as well as statistical and event-driven simulation. CNCL v1.4 (released 01/05/96) is available on the CUG CD-ROM as volume #443. See also the CNCL home page:

<http://www.commets.rwth-aachen.de/doc/cncl.html>





New Products

Industry-Related News & Announcements

ImageFX Introduces FXTools 4.0 Professional Edition

ImageFX has released FXTools 4.0 Professional Edition, a royalty-free toolkit, which adds imaging and multimedia special effects technology to Visual Basic, Visual C++, Delphi 2.0, and other development environments that can host 32-bit ActiveX or 16-bit VBX controls in Windows 3.x, Windows 95, or Windows NT.

FXTools 4.0 contains eight reusable ActiveX and VBX components that give developers control over the display of images, text, and shapes, along with easy playback of sound and video. Each control was designed specifically for interactive development. The sound control supports MCI wave and MIDI audio along with 32-channel sound mixing using DirectSound. The video control supports AVI, MOV, and MPEG file formats. Individual video frames can be displayed with special effects.

FXTools 4.0 costs \$399. Demos can be downloaded from the Internet (<http://www.imagefx.com>) or from CompuServe (GO IMAGEFX).

For more information contact ImageFX, 3021 Brighton-Henrietta TL Rd., Rochester, NY 14623-2749. +1-716-272-8030. FAX: +1-716-272-1873. e-mail: imagefx@imagefx.com. WWW: <http://www.imagefx.com>.

Antares Announces ObjectStar 3.0

Antares Alliance Group has announced the release of Huron ObjectStar 3.0, a cross-platform application development environment for use in the migration of mainframe and desktop legacy applications to three-tiered client/server systems. ObjectStar enables developers to build and distribute mission-critical applications selectively, eliminating the risks of mass migration of legacy applications.

IPT Releases SNIP for Windows

IPT Corporation has introduced SNIP, a C++ programming tool that allows users to specify their code patterns, describe object model design, and generate C++ classes for those design models. SNIP users enter their Data Structure Model and object relationships using the standardized rules for their corporation, group, or team. SNIP can be extended with a company's own code generation templates to capture entire coding strategies that make sense in a company's application area.

The three-tiered model distributes data, application logic, and user interface independently. ObjectStar provides transaction processing monitors, object brokers, and other software needed to implement three-tiered computing for high transaction volumes and for applications supporting thousands of users.

ObjectStar 3.0 starts at \$8,000 per developer seat.

For more information contact Antares Alliance Group, 17304 Preston Rd., Ste. 1200, Dallas, TX 75252. +1-214-447-5500. FAX: +1-214-447-5783.

Thomson Announces Talkinx for Inter-Process Communication

Thomson Software Products has announced the release of Talkinx, a development environment for building applications that require inter-process communication such as client/server, peer-to-peer, and three-tier. Talkinx runs on multiple platforms including UNIX and Windows.

Talkinx supports distributed three-tier applications with a six-function API. The included object library links to the application at build time.

Code results will be obtained automatically with SNIP, because patterns are applied to the code creation process.

SNIP is available for Windows NT, Windows 95, and as a command-line version for UNIX.

For more information contact IPT Corporation, 1076 East Meadow Circle, Palo Alto, CA 94303. 800-944-5468 or +1-415-494-7500. FAX: +1-415-494-2758. e-mail: sales@iptcorp.com. WWW: <http://www.iptweb.com/>

Talkinx generates GUI events asynchronously, eliminating polling by applications. Talkinx also supports other native GUI builders for X and Motif.

Priced at \$2,995, Talkinx is currently available on most UNIX platforms including Sun Solaris, SunOS, HP/UX, SGI Irix, IBM AIX, and OSF/I. Talkinx ships without runtime royalties.

For more information contact Thomson Software Products, 10251 Vista Sorrento Pkwy., Ste. 300, San Diego, CA 92121. 1-619-457-2700. FAX: +1-619-452-2117.

Matra Datavision Announces CAS.CADE v1.3

Matra Datavision has announced the release of CAS.CADE v1.3, an object-oriented software application development platform that provides programmers with a series of reusable software components and a set of predefined object libraries. CAS.CADE includes a class browser that allows the developer to go directly into the data dictionary and navigate graphically to retrieve documentation about the reusable software components. Using cut-and-paste functions, the developer can insert these components

SQA Announces SQA Suite 5 for Windows 95 and NT

SQA, Inc. has announced SQA Suite 5, a testing solution for Windows NT and Windows 95 that includes a scaleable, client/server test repository. The SQA Test Repository integrates the entire testing process and all users of any SQA product on a local or wide area network. Implementing the SQA Test Repository on a large client/server database allows users to implement an enterprise-level client/server testing project. Based on the Sybase SQL Anywhere scaleable database, the SQA Test Repository client/server database offers the portability and recovery architecture of SQL. In

addition, SQA offers a repository implementation option for single-user or small workgroup implementations, which uses the Microsoft Access database.

SQA Suite 5 comes in two versions. SQA Suite: TeamTest Edition, for \$2,995 per user, includes SQA Robot and SQA Manager. SQA Suite: Client/Server Edition, starting at \$12,395, includes SQA Robot, SQA Manager, and SQA LoadTest.

For more information contact SQA, Inc., 10 State St., Woburn, MA 01801. +1-617-932-0110. FAX: +1-617-932-3280. WWW: <http://www.sqa.com>.

where they are needed in the application program. In the CAS.CADE development language, improved memory management reduces the size of code generated to improve engine performance.

For more information contact Matra Datavision, Inc., One Tech Dr., Andover, MA 01810. WWW: <http://www.mdtvus.com>.

Micro Digital Releases smx v3.2

Micro Digital has released their smx v3.2 real-time kernel. smx's main modules are the smx kernel, DOS-compatible file manager, TCP/IP stack, dynamic load module system, smx++ class library, task debugger, protected-mode environment, and windowing utility. These are still sold separately and may be purchased in any combination. One new feature of version 3.2 is an integrated platform that assures any combination of modules will integrate smoothly together.

smx v3.2 supports real mode, 16-bit segmented-protected mode, and 32-bit flat-protected mode. All modules are supported in all three modes. smx also supports the latest versions of 16- and 32-bit Visual C++, Borland C++, Paradigm DEBUG and LOCATE, Soft-Scope, CSILocate, and Periscope.

smx is sold royalty-free for \$3,500. Prices include site development licenses and six months of free support and updates. smx products may be purchased on a 30-day trial basis.

For more information contact Micro Digital Inc., 6402 Tulagi St., Cypress, CA 90630-5630. 800-366-2491 or +1-714-373-6862. FAX: +1-714-891-2363.

Rational Announces Rational Rose for Java

Rational Software Corporation has announced Rational Rose for Java, providing developers the capabilities for object modeling, controlled iterative development, and round-trip engineering. In addition, Rational released a Java-ready extension of the Rational Apex C/C++ environment, enabling software managers to establish and maintain control of their application architecture and providing integration of Java with C and C++ code. Together, the two products offer support for large-scale Java development efforts.

Rational Rose/Java starts at \$2,400 (U.S. price) for the Windows version, which includes the first year of support. Rational Apex C/C++ pricing ranges from \$7,500 to \$12,500, depending on configuration and bundling options.

For more information contact Rational Software Corporation, 2800 San Tomas Expressway, Santa Clara, CA 95051-0951. 800-728-1212 or +1-408-496-3600. FAX: +1-408-496-3636. e-mail: product_info@rational.com. WWW: <http://www.rational.com>.

REQUISITE Announces RequisitePro v1.0

REQUISITE, Inc. has announced RequisitePro 1.0, a team-based requirements management tool that integrates traceability and the version control product, PVCS Version Manager. The Windows-based tool provides the ability to follow traceability relationships through a structured hierarchy. For example, when a high-level user needs a change, a developer can quickly identify which software elements must be altered. A tester can pinpoint which test protocols must be revised. Managers can better determine the potential costs and difficulties of implementing the changes. In addition, RequisitePro builds a requirements database that is attached and synchronized to the projects specifications and test plans. This repository holds a development team's documents, requirements, and traceability relationships.

With RequisitePro, individual requirements are itemized from within Microsoft Word project documents and managed with an integrated database. All information is centralized in one shared repository. Requirements documentation, manage-

DataViews Corporation has announced DV-Centro 1.2, a development tool for building Visual Programming Languages (VPLs) under Windows NT. DV-Centro allows programmers to create true VPL systems based on custom symbols, shapes, and objects connected to data. Based on an object-oriented C++ framework, DV-Centro features a suite of graphical tools that allow users to create reusable components for complex VPL applications. Also featured is an Interactive Control Engine that manages events between the user display and the application components.

DV-Centro 1.2 features a set of coordinated C++ classes optimized for managing the relationships between graphics and data in VPL applications. This framework decreases the amount of code that has to be programmed, tested, and debugged. DV-Centro also includes a set of C++ utilities for memory management, infinite undo/redo, object storage/retrieval, and journaling.

DV-Centro 1.2 supports Windows NT 3.51 with Microsoft's Visual C++ compiler, HP-UX 10.x, and Sun Solaris 2.4 with a native C++ compiler. Development seats for DV-Centro begin at \$15,000.

For more information contact DataViews Corporation, 47 Pleasant St., Northampton, MA 01060. +1-413-586-4144. FAX: +1-413-586-3805. e-mail: info@dvcorp.com. WWW: <http://www.dvcorp.com>

ment, traceability, and reporting are now performed with one tool.

For more information contact REQUISITE, Inc., 4720 Table Mesa Dr., Ste. F, Boulder, CO 80303. 800-732-4047 or +1-303-499-9177. FAX: +1-303-499-9535. e-mail: 74067.612@compuserve.com.



NetManage Ships Swift 2.0 IntraNet Access Package

NetManage, Inc. has released Swift 2.0 and Swift IntraNet Package 2.0. Swift provides company-wide host access services over multiple protocols for Windows 3.1, Windows 95, and Windows NT. Based on the host access components of Chameleon IntraNet Desktop suite, Swift 2.0 supports access to IBM mainframes, AS/400 and VMS/UNIX hosts over TCP/IP, SNA, and asynchronous connections. It includes host emulation services, NetManages' WebSurfer HTML browser, industry-standard SMTP e-mail, and the NS/Router client software.

Swift 2.0 costs \$199. The Swift IntraNet Package is available in a minimum quantity purchase of 20 copies for \$420 per copy.

For more information contact NetManage, Inc., 10725 North De Anza Blvd., Cupertino, CA 95014. +1-408-973-7171. FAX: +1-408-257-8780. e-mail: marketing@netmanage.com. WWW: <http://www.netmanage.com>.



QSI Ships ProTEST 3.5 for Windows

QSI has begun shipping ProTEST 3.5, a test process management system designed to assist in the development of a formal and structured test methodology, and manage the testing process. Enhancements include the ability to assign weight factors to a test case procedure, test case summary reports that provide detailed status of the test case, and an online "Regression Testing - To Do" file that allows direct access to the test cases that have failed, in order to include them as part of the Regression Test Phase of the project. Additional enhancements include test case status update capabilities that allow testers to pick and choose which test cases to update with a Pass/Fail or Incomplete status. Another enhancement is a Quick View function that provides a display of all the test case procedures in



ObjectSpace Announces Web<ToolKit>

ObjectSpace, Inc. has announced Web<ToolKit>, an ANSI/ISO-compatible C++ class library for producing HTML World Wide Web pages. Web<ToolKit> supports page creation using a set of C++ classes representing HTML elements, including text, links, graphics, tables, forms, frames, and widgets.

Web<ToolKit> uses ANSI/ISO standard containers and strings. Version 1.0 operates with existing ObjectSpace products, including STL<ToolKit> and Systems<ToolKit>. The product can also be used with Borland 5 STL and the GNU STL.

Cost for the software is \$349 for PC platforms and \$475 for UNIX platforms. Special bundled pricing is available with STL<ToolKit>, ObjectSpace's implementation of the Standard Template Library, and Systems<ToolKit>, ObjectSpace's second generation C++ toolkit.

For more information contact ObjectSpace, Inc., 14881 Quorum Dr., Ste. 400, Dallas, TX 75240 +1-214-934-2496. FAX: +1-214-663-9099. e-mail: training@objectspace.com. WWW: <http://www.objectspace.com>.

one screen, along with the associated requirements, problem report numbers, description of the procedures, and the status of each individual test case procedure.

For more information contact Quality Systems International, Inc., 416-420 Highland Ave., Cheshire, CT 06410. 800-557-9787 or +1-203-699-9787. FAX: +1-203-699-9789.



Programmer's Paradise Releases Communications Paradise Catalog

Programmer's Paradise, Inc. has released a new catalog featuring products for Internet and intranet development, communications, and connectivity. Published quarterly, the Communications Paradise Catalog features HTML browsers, Java applet creation tools, web servers, authoring and editing tools, search engines, and an array of communications and connectivity products including modems and peripherals, navigators, TCP/IP, and e-mail software.

For more information contact Programmer's Paradise, Inc., 1163 Shrewsbury Ave., Shrewsbury, NJ 07702-4321. +1-908-389-8950. FAX: +1-908-389-9227.



IBM Expands VisualAge for C++ to Cover Windows

IBM has added support for Windows NT and Windows 95 to VisualAge for C++. VisualAge for C++ for Windows provides an integrated C++ development environment that includes a data access builder, which automatically generates objects and parts for accessing relational data.

The data access builder has been enhanced in the VisualAge for C++ for Windows release to support Open Database Connectivity (ODBC). In addition to native support for DB2 for Windows NT, this builder can access Oracle, Sybase, and other relational databases supported by ODBC.

VisualAge for C++ for Windows also includes an optimizing compiler that supports the Pentium ProChip, a syntax-highlighting editor, class browser, program debugger, and performance analyzer.

VisualAge for C++ for Windows v3.5 costs \$449. Customers using any competitive product with a C++ compiler on Windows will be offered a special upgrade base price of \$299.

For more information contact International Business Machine Corporation, Route 100, P.O. Box 100, Somers, NY 10589. 800-426-3333. WWW: <http://www.software.ibm.com/info/pr0305.html>.



Micro Focus Ships Revolve Year 2000 Add-On

Micro Focus has announced the shipment of the Micro Focus Revolve Year 2000 Add-On and Challenge 2000 packages. The Revolve Year 2000 Add-On is one component of Micro Focus Challenge 2000 v2.0, a product that includes tools and services for addressing all phases of the millennium date-change problem, from assessment and identification through implementation and testing.

The Revolve Year 2000 Add-On tracks the analysis and changes required for Year

Mainsoft Ships MainWin Test 3.0

Mainsoft Corporation has begun shipping MainWin Test 3.0, an automated testing tool compatible with Microsoft Test 3.0 for testing Windows applications ported to UNIX platforms. MainWin Test 3.0 offers developers a testing library featuring tests for the latest Windows controls; powerful and easy-to-learn TestBasic scripting language; an integrated development environment with visual syntax editor, compiler, and debugger; and screen capture/compare and dialog capture/compare utilities.

MainWin Test for UNIX offers compatibility with Microsoft Test for the PC,

both in test-script development and execution. MainWin Test allows the reuse of the same testing scripts created for validating the Windows product.

MainWin Test will be supported on DEC, HP, IBM, SCO, SGI, and Sun UNIX. Developer pricing for MainWin Test starts at \$5,000.

For more information contact Mainsoft Corporation, 1270 Oakmead Parkway, Ste. 310, Sunnyvale, CA 94086. +1-408-774-3400. FAX: +1-408-774-3404. e-mail: info@mainsoft.com. WWW: <http://www.mainsoft.com>.

2000 compliance using assessment, analysis, and modification features. The Year 2000 Add-On enables users to examine, visualize, and assess systems consisting of thousands of programs and millions of lines of code.

Micro Focus is also introducing two customized packages designed to match the needs of organizations addressing their Year 2000 problem. These packages, the Challenge 2000 Pilot and the Challenge 2000 Compile and Test Option, provide a Year 2000 team with the tools, training, and support to launch their assessment, analysis, and modification efforts. The Challenge 2000 Pilot costs \$46,500. The Challenge 2000 Compile and Test Option pricing starts at \$43,000.

For more information contact Micro Focus, 2465 East Bayshore Rd., Palo Alto, CA 94303. +1-415-856-4161. FAX: +1-415-856-6134. WWW: <http://www.microfocus.com>.

PLATINUM Releases Paradigm Plus v8.0

PLATINUM technology, inc. has released Paradigm Plus v8.0, an object-oriented analysis and design tool that supports version 0.8 of the Booch/Rumbaugh/Jacobson Unified Modeling Language (UML), Enterprise Component Modeling (ECM), and round-trip engineering. With Paradigm Plus, companies create "components" that divide information systems into manageable sections. Components are stored in the Paradigm Plus object repository where they are organized and synchronized with source code and documentation.

Paradigm Plus generates code in C++, Ada, Smalltalk, PowerBuilder, Java, ObjectPro, and Forte, as well as RDBMS and ODBMS schema definitions. It also reverse engineers C++, Smalltalk, PowerBuilder, ObjectPro, and Forte code. The combined forward and reverse engineering capabilities of Paradigm Plus, without any loss of data, provide round-trip engineering, which enables iterative development. Utilizing a distributed object repository and open architecture, Paradigm Plus optimizes reuse by allowing access to all stored objects and relationships in a multi-user environment. Paradigm Plus components also can be published on the Web as HTML documents, where they can be viewed, downloaded, and reused.

Paradigm Plus 3.01 is available for Windows 3.1, Windows NT, and Sun Solaris.

For more information contact PLATINUM technology, inc., 1815 South Meyers Rd., Oakbrook Terrace, IL 60181-5241. 800-442-6861 or +1-708-620-5000. FAX: +1-708-691-0710. e-mail: info@platinum.com. WWW: <http://www.platinum.com>.

JBA Introduces JBA Guidelines v3.3

JBA International has introduced JBA Guidelines v3.3, a software development environment for OS/2 WARP and OS/2 v2.1, for generating GUI-based client/server applications. Guidelines has incorporated a messaging layer that ensures that all applications built today are object-enabled. Version 3.3 enhances the distributed message layer by providing an interface

between standardized object calls and the code servicing these calls. This code can be a mixture of existing logic and more advanced CORBA-compliant objects.

Currently Guidelines is packaged in two versions: Guidelines for Desktop Developers, with prices starting at \$595, and Guidelines for Corporate Developers, priced at \$7750 per seat. A Base Pack, in essence a Guidelines tester, is free and can be downloaded from CompuServe (GO Guidelines) or the Internet (<http://www.jba.co.uk>).

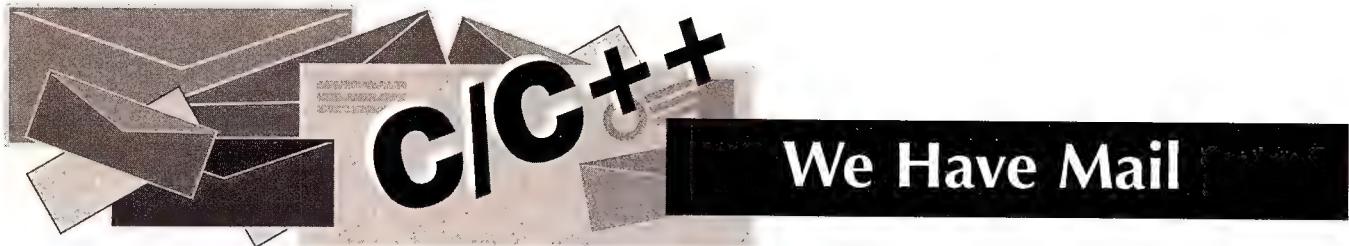
For more information contact JBA International, 161 Gaither Dr., Ste. 200, Mt. Laurel, NJ 08054. 800-522-4685 ext. 3074. FAX: +1-609-222-1952. WWW: <http://www.jba.co.uk>.

Art Technology Group Introduces Dynamo 2

Art Technology Group has introduced Dynamo 2, a Java-based, Internet application engine. Dynamo 2 enables Web developers to prototype, build, and deploy scalable, high-volume, and database-driven Web applications requiring dynamic page generation and personalized content. Based on open standards, Dynamo 2 allows Web developers to embed native Java code (not Javascript) into HTML files to extend functionality on the server side. Java code is compiled on the server to customize everything from the visual characteristics of the page to behaviors of traditional client-side Java applets.

Pricing for Dynamo 2 is \$5,000 per license, and one year of support can be purchased at \$750/year per license. Support is available via the Web, where users can get up-to-date technical information, receive technical support, and download new "Gears" and releases.

For more information contact Art Technology Group, +1-617-859-1212. e-mail: info@atg_dynamo.com. WWW: <http://www.atg-dynamo.com>.



Letters to the editor may be sent via email to cujed@mf.com, or via the postal service to Letters to the Editor, C/C++ Users Journal, 1601 W. 23rd St., Ste 200, Lawrence, KS 66046-2700.

Dear Dr. Plauger,

I am sending you the following suggestion in the hope that you can communicate it to the ANSI C++ committee. I realize they are close to freezing the language, and may already have done so by now, but still I thought it was worth a try.

A widely acknowledged shortcoming of C++ is the visibility (in header files) of private attributes and methods. When a new private method is added, files that include the header only for the public interface are needlessly recompiled. This can become a major nuisance in complex projects with many interdependent files.

I propose a simple addition to the language to fix this. It would not affect existing code. The idea is to allow private member functions to be defined in the implementation file without being declared in the class declaration. Such functions would have file scope, just like static functions commonly used in C. Unlike static functions in C, however, they would have access to all the private attributes and methods of the class to which they belong. The syntax could use the keyword `static`, thus:

```
static void MyClass::  
    method_with_file_scope( void ) {  
    // function defined here  
}
```

The compiler would disallow the declaration of the same function in the header file. If one more meaning of the already overloaded keyword `static` is objectionable, the keyword `private` (or some new keyword) could be used instead. Like conventional private member functions, these functions would not be accessible to a derived class. They would not be accessible to a friend class unless the friend is implemented in the same file.

In my experience, private member functions change more often than data members and other functions, especially during development, and they usually all live in the same source file. Although this simple change would not eliminate all needless recompliations, it would surely help.

Sincerely yours,

Jack Wathey

Dear P.J. Plauger,

Declaration of non-virtual private member functions or, did Bjarne nod? Having been a C programmer for many years, I was looking forward to the opportunity of doing some serious work in C++. But I soon found vexation and frustration.

I'm used to hiding the implementation for an "object" in C by declaring what you need to know to use it in a header file, and the rest in a private header file or in a .c source file. But it turns out that because member functions have to be declared in the header, a lot of implementation shows. I'm talking about the sort of functions you would usually declare as static. Like

```
for ( i = 0; i < something ; i++) {  
    dosomethingto(fred[i]) ;  
}
```

where `dosomethingto` is a strictly local function invented on the spur of the moment. But in C++, you have to declare it in the header. If I try this, my current compiler says something like "`dosomethingto` is not a member function of class `C`." This means that you have to advertise your implementation to all your users in the header file.

In the ARM, this requirement is justified to prevent people from "hijacking" the class, like this:

```
class C {  
private:  
    int t;  
...  
}  
void C::hijack() { t++ ; }
```

But I think that this goes against Stroustrup's admirable attitude generally, that protection is there to prevent you inadvertently doing something you didn't mean to do, and not to stop you from doing what you want to do. If you wanted to hijack the class even more thoroughly you could write one or all of the member functions yourself. With most compiler/linker combinations, you would get your own implementation and no warning that there was another in the library. But this isn't called "hijacking," it's called "implementing a class!"

My proposition is that a declaration/definition of a member function that is not declared in the class should be accepted, it should be deemed to be private, its name should have file scope. Being private it would only be able to be called from other member functions (that takes care of the hijacking issue). One good thing about this is that it is not an addition to the language, and it doesn't break any existing correct code. Such functions don't affect the size of the members of the class, so it's feasible (as far as I can see).

Maybe this idea has already been discussed, and maybe even proposed. Maybe it's even on its way! Meanwhile, I have adopted another approach for my own implementation, which would go like this for class `C`:

```
class C_imp; // implementation of C  
class C {  
    friend class C_imp;  
private:  
    C_imp * cp ;  
public:  
    // public functions of C  
    ...  
protected:  
    // functions derived classes can  
    // inherit  
    ...  
};
```

So all the private data and functions of this class can be declared in my private header file `<c_imp>`

The constructor for `C` has to do a new on the `cp` member — the only private member. The functions of `C` then provide a front for the real functions in `C_imp`. It also hides what possible classes I may or may not have inherited from in building my class.

I agree that this is a bit of a kludge, but I think, for a class of any size, it's pretty important to hide the implementation. I still have to stick declarations for all my tiny functions into the (local) declaration of `C_imp`, but I don't have to burden the customer with them. What's more, when they get a re-release of the code, they don't have to recompile, but only re-link.

I would be interested to know what you think about this proposal.

Yours,

Dan Oestreicher

The scheme you outline is a technique sometimes used to provide additional hiding of implementation details. Even it can be hijacked, by someone who knows, or can guess, enough of those hidden details. "Private" doesn't necessarily mean "secret." It is more a way to minimize accidental use of internal information about a class. I know of no proposals to add the kind of extension that both you and Mr. Wathey suggest, and I doubt they'll be forthcoming at this late stage. — pjp

Dear Mr. Plauger,

I am working for a software company which has the majority of the products written in C. We think that we should shift to C++ where possible. We are trying to gather information about how it can be done from those who experienced the transition. What is a good way to start? Is there a formal procedure? Where can we get information about this? Any advice is appreciated. Thanks.

Mustafa Yorgancioglu
Tempe, AZ

A number of articles have been written in the trade press, over the past few years, describing experiences in moving from C to C++. You will find a few in back issues of CUJ. In fact, Bobby Schmidt's column, "The Learning C/C++-urve," focuses especially on how to migrate code from C to C++. The basic rule is to take small steps, preserving old code in a mixed environment as much as possible. C++ works best for writing new code designed from the outset with object-oriented design principles in mind. — pjp

Dear Mr. Plauger:

You might as well do it. Get it over with. Do us all a favor. Please go ahead and change the name of your once fine magazine to *The C++ Users Journal* so it accurately reflects the content within, and so I can no longer accuse you of false advertising. I for one still refuse to give up to the legions of bloated C++ programmers spawned by Mr. Gates and his like. Long live C for embedded programming...

Steadfastly resisting the hype,

Mark S. Outson

I think it's the programs that are bloated, not the programmers, and I have trouble blaming Bill Gates for C++. Still, we have been running light on "strictly C" articles in recent months. You'll be happy to know we've taken steps to acquire more C articles. Since C underlies so much of C++, a good C article can be relevant to all programmers — it doesn't matter whether your forte is C or C++. Of course, as the next letter demonstrates, not everyone shares your perception of our content. — pjp

Dear Mr. Plauger,

Alas, I find I must bid a fond farewell to *The Journal*, and it does seem fair, considering the many letters you receive extolling the virtues of good old CUJ, that I should address this to you. After all, balance is important. Right?

I, unfortunately, am one of those who've found that with each passing issue less and less of any worth (program-wise) appears in *The Journal*. Come to think of it, CUJ has never been long on programming information.

Lots of coding, to be sure, but except for an occasional article on how to handle this or that side-saddle piece of equipment, or how to rework a few obscure statements in an algorithm of rare vintage and even rarer utility, the overriding gist of what's published in *The Journal* (its one claim to fame) has always been:

We peddle esoterica and are loudly proud of our total commitment to the protection and preservation of the purity of the cryptic qualities of the work performed by ANSI C Standards Committee XYZ789.

Hey, isn't that how Gurus earn a living?
Truly,

Mark Rhyner

I personally earn most of my income writing code in C and C++ that people find worth licensing. But we can't please everyone with this magazine. Sorry to see you go. — pjp

Dear Mr. Plauger,

I am a novice programmer in C. The past two months I've been trying to teach myself the language. The instrument that I've been using

X/Motif Widgets UNIX/VMS

- AtPlotter Widgets: line, bar and XY charts, text, flexible programming, online processing, print, C source
- BXWidgets: portable TextTable, Center TableManager and Hypertext Help Widgets for X/Motif, C source
- XView: portable viewer Widgets and Viewer for technical formats like TIFF, Formtek, PCX, GIF, MI, HPGL, MI, CALS, C alComp, Medusa. All formats with zoom, pan, sheet, redraw, undo, rotate. TIFF also with crop and redline.
- GLIB: 2D graphics library for X/Motif and MS Windows. Supports picking, multiviewing, world coordinate system, primitives like ellipse, arc, chord, curve, bezier, spline, bitmaps, fill, text.

Whalen Soft. (USA) ☎ +1 815 729 3130
F.Ritz Systems (D) ☎ +49+7731977030

□ Request Reader Service #193 □

C/C++ Professionals

Manpower Technical, a world leader in the Technical Services industry, has current and upcoming contract openings for experienced C/C++ personnel with expertise in any of the following: Software Development, UNIX, OOA, OOD, MOTIF, Telephony, Telecommunications, ADA, Device Drivers, GUI.

Positions are located throughout North America.

Interested persons are encouraged to send their resume to:

Manpower Technical
Dept. 32A
P.O. Box 2053
Milwaukee, WI 53201

Phone: 800-558-6992
FAX: 414-332-0378

□ Request Reader Service #182 □

Dr. DeeBee® ODBC Tools

Tools for ODBC development

16 &
32 Bit

Ever wonder why
your ODBC app is not
working? Why it's just
too slow? If the ODBC
driver is OK?

Dr. DeeBee utilities reveal the inner workings of ODBC.

Dr. DeeBee ODBC Driver Kit

NEW VERSION! Connect proprietary databases to Access, Visual Basic, and PowerBuilder with our Dr. DeeBee ODBC Driver Kit

SYWARE Inc.

P.O. Box 91 Kendall, Cambridge, MA 02142
617-497-1376 Fax 617-497-8729
<http://www.syware.com>

□ Request Reader Service #200 □

Software Consultants

Want to Make More Money?

Why do some software consultants make \$200K a year? Simple, they read *The Software Consultant*, a monthly newsletter. Here's what's in every issue:

- **Increasing Sales** - Learn red hot business tricks guaranteed to increase your profits. Things like money making advertising strategies, easy ways to find clients and super ways to close big deals.
- **Trends** - We'll show you what the big money consulting jobs are and how to get your share of them. For example, why doctors are hiring lots of consultants (*hint: it's not web pages*) and how you can cash in on the trend.
- **Systems** - Learn easy ways to tackle tough problems like testing, integration, and reducing risks.
- **How To** - Learn how professionals get requirements, finalize designs and sell off completed projects.
- **Economics** - Know where the economy is headed and what that means to your consulting business.

→ Write for free information today.
No obligation of course.

Reo Centered Publishing, Dept. C1,
330 W. Dorado Circle, Litchfield Park, AZ,
85340 e-mail Reocenter@aol.com

□ Request Reader Service #191 □

The Complete C++ Math Library

M++ Version 5.2

A full multidimensional array language
Complete LINPACK and EISPACK
Full suite of signal processing functions
Support for UNIX, DOS, 95, NT, LINUX
Support for Persistent & Virtual arrays

Advanced optional modules

Statistical Utilities	Optimization
Visualization	Numerical Integration
Financial Utilities	Ordinary Diff Eqs
Least Squares	Finite Element Modeling

Free demo on website

<http://www.wolfsnet.com/~dyad>

Dyad Software Corporation

email dyad@wolfsnet.com

(206) 637-9426 V (206) 637-9428 F

□ Request Reader Service #175 □

C/C++ Users Journal — September 1996

Little Big LAN

Costs just \$75 per LAN!

- * Connect via Ethernet, Arcnet,
- * via serial, parallel, or modem
- * Low memory use - about 50k
- * Dos 6, and Windows compatible
- * Link up to 250 nodes, still \$75
- * Share almost anything
- * Netbios and file/record locking
- * Version 2.0b

Skeptical? We make believers!



Information Modes
P.O. Drawer F
Denton, TX 76202
817-382-7407 Fax
1-800-628-7992

□ Request Reader Service #179 □

Data Structures

The MemSL™ is the most complete Data Structures Library for C & C++.

Includes:

Single and Double Linked lists,
Multi-Dimensional Arrays, Hashing Tables,
Binary Trees, Balanced Trees, Queues,
Deques, Priority Queues,
AVL and Threaded AVL Trees,
Memory Tracing and Debugging,
And More for only \$79.99!
Source Code Included, Royalty Free.

Windbase® Software Inc.

P.O. Box 10115 • Glendale, AZ 85318-0115

602-561-8788

Fax: 602-561-6490 BBS 602-780-9692, N81

□ Request Reader Service #205 □

OPT-TECH SORT/MERGE

★ New - Version 5

High performance Sort/Merge>Select utility. Run as a stand alone utility or CALL as a subroutine.

Supports most languages and filetypes including Btrieve and dBase. Unlimited filesizes multiple keys and much more.

**MS-DOS, Windows \$149
Win95, OS/2, UNIX \$249**

Call to order or for free info.

Opt-Tech Data Processing

P.O. Box 678

Zephyr Cove, NV 89448

(702) 588-3737

□ Request Reader Service #187 □

CAREER ADVANCEMENT

Let C Associates Expand Your Technical Career

Placement Specialists In:

C & C++ Programming, UNIX System Administration, ORACLE, Sybase, Informix, Progress, FoxPro, other relational databases.

Network Design & Management, Sun Workstations, Graphics Development, MS & X Windows Applications

Please mail or fax your resume to:

C C Associates
ATTN: John Capozzi
P.O. Box 15420
Washington, DC 20003-0420

Phone: 202-544-0821 Fax: 202-547-8357

Do it today!

□ Request Reader Service #170 □

Great Software Jobs, Nationwide!

Skilled software developers like you are in demand all over the country. TravelTech Staffing can place you on a six to twelve month contract, or regular full-time position just about anywhere you want to go. We offer challenging positions for serious professionals.

When you accept a contract position through TravelTech, you can customize a benefits package to suit your needs. Choose from an attractive menu of benefits including top-pay, free round-trip transportation, liberally subsidized high-quality housing, health and dental insurance and more. It's up to you!

To learn about our program, call the toll-free number below for a free information package. Be sure to fax or e-mail your resume so one of our technical recruiters can be looking for the perfect position for you. More information, including immediate openings, can be found on our website.

1-800-282-0360

E-mail: resume@traveltechstaffing.com
Website: www.traveltechstaffing.com
Fax: 1-800-792-7996

TRAVELTECH
STAFFING, INC.

Cash in on the Voice Bonanza!

Get your piece of the \$3 billion Voice Industry - develop your own Computer Telephony Apps!

Add voice power to your Windows apps with **TIF DLL™** (CT Mag's Product of the Year), our popular interface to Dialogic multi-line telephony hardware.

Audio ToolBox™ lets you batch process and convert audio files to & from standard formats. Adjust volume, filter, trim silence, more! Add to your apps with **ToolBox SDK**.

Call and ask about our **VFEedit™** professional voice prompt editor, **Scribe™** transcription utility, **VoxFonts™** text-to-speech library, plus more great products!

ISI 800-469-4874
VISA, 2118 Wilshire Blvd, #973, Santa Monica, CA 90403
310-392-6780 Fax: 800-234-FXIT / 310-392-5511 sales@vinfo.com
BBS: 310-392-6610 www.vinfo.com (Download Working Demos)

□ Request Reader Service #202 □

Need Precision Timing Under Windows?

ExacTicks (\$129.95) provides precise timing, alarms, delays, and event scheduling under the Windows 16 and 32 bit x86 variants. **ExacTicks** is the definitive Windows timing library for execution profiling, data acquisition, and process control applications. Includes 16 & 32 bit DLLs, DLL interfaces for C, C++, Delphi, and Visual Basic. Full source code and DLL distribution license included.

Using Borland's BGI Graphics Library?

BGI SVGA Video Driver Toolkit (\$129.95) supports up to 1280x1024 in 16 & 256 colors, PCX/BMP support, SVGA mouse, 24 different SVGA chipsets + VESA.

BGI Printer Driver Toolkit (\$129.95) provides b/w and color printing via BGI on nearly every printer and plotter on the planet, all at the device's full resolution - not a screen dump!

BGI Font Toolkit (\$89.95) fixes BGI text bugs, adds arbitrary rotation, underlining, bolding, italicizing, convert and use TrueType fonts with our supplied converter.

All BGI toolkits include full source, Borland C/C++/Pascal real & DPMI16 support, driver &.obj distribution license.

VISA, MC, AMEX. Add \$5 shipping USA, \$10 elsewhere.
30 day "No Questions Asked" return policy.

Ryle Design PO Box 22, Mt. Pleasant MI 48804 USA
Voice: 517.773.0587 73047.1765@compuserve.com
Fax: 517.775.5530 http://www.sojourn.com/~ryle/web/

□ Request Reader Service #195 □

The Weekly MFC EXTension

- The only MFC extension class library that delivers new class(es) each week
- Updated through Internet e-mail
- Includes all classes available upon subscription
- Free sample classes available as demo
- Only \$249 annual subscription
- Classes available include : Smart pointers and brilliant groups: GRID; MAPI; Dockable Property Sheet; JPEG compressor & decompressor; Extended Statusbar; OLE automation Server; Structured Exception Handling; Client/Server communication(includes JAVA client) and many, many more ...

Now is the time to hire your most productive developer at \$249 a year !!

Contact or visit Peripherie Phone: 303-682-5297
eMail: info@peripherie.com Fax: 303-772-9430
http://www.peripherie.com

□ Request Reader Service #189 □

Can Parsing Be Fun and Easy?

Parsing input used to mean tricky logic with lots of places for bugs to hide, slipped schedules, difficult maintenance.

No more! Now you just describe your input and **AnaGram** writes your parser. You get flexible, reliable, maintainable parsers on time and under budget.

Validate user input fields, interpret data base queries, read configuration files, interpret macros, or write a compiler.

Software developers around the world are turning to **AnaGram**. Shouldn't you too? Call now for your free trial copy.

AnaGram™ by **Parsifal Software**
P.O. Box 219, Wayland, MA 01778

1-800-879-2577 Voice/Fax 1-508-358-2564
CIS:72603,1763 jholland@world.std.com

□ Request Reader Service #188 □

has been the Waite Group's beginning primer book, *The New C Primer Plus*. My question to you is the following: Are the articles presented in your periodical *C/C++ Users Journal* suitable for an individual with my skill level in C programming? Secondly, if my skills are not adequate what materials can you suggest that I investigate and acquire to raise my skill level adequately so as to take advantage of the material presented with in *The C/C++ Users Journal*?

Any advice or encouragement would be greatly appreciated.

Thank you for your time and attention.

David Hodge
dlhodge@netcom.com

I suspect that only a small fraction of our articles will make sense to you right now, but that should change rapidly as you gain experience with C. The only way I know to learn any programming language is to practice writing code, preferably useful code. The more you write, the more you learn. Good luck. — pjp

Hi Mr. Plauger,

A few months ago I started programming in C and Visual C++ again. I used C in some of my number crunching programs a few years ago. I am currently working on an embedded control system on the iRMX OS and multibus II arch (a real-time program).

I came across your magazine a few months ago...very interesting stuff. I purchased the May '96 and several of the '95 publications, all that I could find. I am in the process of ordering the CD.

Question: Has your magazine ever published any articles in the following areas: database and MPEG?

BTW I was going through the Q&A in one of the '95 pubs and wow...I was able to solve one of my lvalue/rvalue errors...many thanks!! Keep up the good work.

Cheers

Anand

The May 1996 issue which you purchased includes a database-intensive article, "Object Persistence with Relational Databases," by Chuck Allison. Another recent database-related article is "C Database Programming with ODBC," by Alex Ragen, in the November 1995 issue. We have yet to run an article on MPEG, but we would not rule it out as a future topic.

We're glad you found Pete Becker's Q&A column so immediately and tangibly useful. I imagine many of our readers have had the same experience. Of course, if you want to be sure you don't miss out on anything in the future, you could always subscribe by sending an e-mail to rorders@mfi.com (hint, hint). Thanks for reading CUJ and thanks for writing. — mb

Dear Mr. Plauger

I would like to open with a compliment on your articles in *C/C++ Users Journal*; they are both informative and well written.

I have been seeking a solution to a C++ problem for several weeks now and am no closer. I had thought to write to you earlier but did not wish to waste your time on a simple problem. For me it no longer seems simple.

**C/C++
Programmers
Market**

C and C++ DOCUMENTATION

!! VERSION 6.0 !!

- **C-CALL (\$69)** Graphic-tree of caller/called functions, cross-ref, file/function index.
- **C-CMT (\$69)** Creates/inserts/comment-blocks for each function, listing the functions and identifiers used by it.
- **C-METRIC (\$59)** Counts path complexity, counts comments, code, 'C' statements.
- **C-LIST (\$69)** Lists and action-diagrams, or reformats into standard formats.
- **C-REF (\$69)** Creates cross-reference of local/global/define/parameter identifiers.
- **C-DOC (\$199) Package** All 5 programs integrated as DOS program. <10,000 lines. V6.0 C-BROWSE Windows graphic-tree viewer.
- **C-DOC Professional (\$299)** DOS, Windows, OS/2. 3-ring binder/case. <1,000,000 lines.

30-DAY Money-back guarantee. CALL NOW!

SOFTWARE BLACKSMITHS INC.

6064 St Ives Way, Mississauga Voice/Fax (905) 858-4466
ONT Canada L5N-4M1 <http://swbs.idirect.com>

Please see Ad Index for our larger ad.

□ Request Reader Service #199 □

TRANSLATORS

**ASM360/370 ASM/86 MASM
Motorola ASM680X ASM680X0
COBOL FORTRAN PASCAL**

Intel PL/M X86 PL/I

- Translate legacy applications to "C"
- Versions for PLM51, 80, 86, 96, 286, 386
- Versions for most Cobol & PL/I dialects
- Generates "C" source and listing files
- Protected mode option for large files
- Easy to use; User's Guide; for MS-DOS
- From \$575 ea. plus s/h; update service
- Translation services & custom dialects

Micro-Processor Services, Inc.
92 Stonehurst La., Dix Hills, N.Y. 11746
(516) 499 - 4461; Fax: (516) 499 - 4727
Email: mpsinc@netusa.net
WWW page: <http://www.mpsinc.com>

□ Request Reader Service #184 □

16 & 32 bit Serial Comm Libs

Personal Communications Library DOS (PCL4C) or Windows (PCL4W) communications library for C/C++. Run 20 ports to 115200 baud. Supports 16 & 32 bit DPMI, multiport dumb cards such as Digiboard & BOCA. Use any IRO & UART address. Interrupt driven, with HW flow control, 16550 UART, and modem AT command support. 38 functions. Source code included. \$75 + s&h

Personal Protocol Library DOS (PPL4C) ASCII, XMODEM, YMODEM and ZMODEM protocols. Source code is included. With script compiler and interpreter. Requires PCL4C above. \$40 + s&h

Includes printed manuals & one year support. Get FREE shareware version from our BBS (14.4 KB). FTP: [ftp_marshallsoft.com](ftp://marshallsoft.com). See review in Windows/DOS Developers Journal (March 95 issue)

MarshallSoft Computing, Inc.

Post Office Box 4543 Huntsville, AL 35815
(205) 881-4630 Voice, 880-0925 FAX, 880-9748 BBS
email: info@marshallsoft.com
web: www.marshallsoft.com

□ Request Reader Service #183 □

Employment Service

SALARIED SOFTWARE DEVELOPMENT AND SUPPORT ENGINEERS

Clients and affiliates nationally...
Clients pay our fees (always), and
your interview and relocation
expenses (usually).

RSVP SERVICES

trusted by computer professionals
since 1966

Ref: UJ

PO Box 8369 Cherry Hill NJ 08002-0369
Voice: 800/222-0153
Fax: 609/667-2606

Internet: npa1621@connectinc.com

mail/fax/email resumé

□ Request Reader Service #194 □



Volt Computer
Services

8461 154th Ave. NE
Redmond, Wa. 98052
Phone:
(206)702-9000

Fax:
(206)702-0315

Email:
jef_vileta@msn.com

VOLT

**SOFTWARE
DEVELOPERS**
Immediate openings
on Redmond,
Washington Campus

- ▼ C, C++
- ▼ Interactive Television
- ▼ Multi-Media
- ▼ Internet

Volt offers competitive wages as well as a comprehensive benefits package

Request Reader Service #203

STOP WASTING TIME CREATING LIST BOXES!

ListBox Calculator™

ListBox Calculator™ eliminates the tedious guesswork of setting up list boxes, combo boxes, and column headings. Easy to use for beginning through advanced programmers. ListBox Calculator automatically determines

- List/combo box width
- List/combo box tab stops
- Column heading sizes and positions

For all languages and most popular fonts

\$49

Visa, MasterCard

Sheet Bend Software

4061 E. Castro Valley Blvd., #197
Castro Valley, CA 94552

Phone: 510 581-2671 Fax: 510 581-8350

Request Reader Service #197

Developer Jobs!

Internet: das@scientific.com

Commercial software developers should register with Scientific Placement. R&D jobs for software engineers, SQA, product managers, etc. Nationwide contacts with both large and small companies including start-ups. Many clients develop and publish commercial software products. Most develop for Windows, NT, Macintosh, OS/2, and Unix based platforms. We also recruit in other leading edge technology areas such as PDA, low level and real-time, compilers, etc. Managed by graduate engineers. Send résumé or call to get an assessment of your marketability. Never a fee.

Scientific Placement, Inc.

800-231-5920 Fax 800-757-9003
<http://www.scientific.com>

Compuserve: 71250,3001 AOL:davessmell
CJ, Box 19949, Houston, TX 77224
713-496-6100 Fax 713-496-0373
CJ, Box 71, San Ramon, CA 94583
510-733-6168 Beth@spica.bdt.com
CJ, Box 202676, Austin, TX 78720-2676
512-260-0123 lej@zilker.net



Request Reader Service #196

C/C++ Users Journal — September 1996

Market and Publish Software!

WishBone Publishing, Inc. seeks software products to publish in the marketplace. If you have developed an innovative and creative game, program or utility, WishBone will review your program, determine its market potential, and publish the finished product.

WishBone
PUBLISHING, INC.

1101 17th Street, N.W., Suite 408
Washington, D.C. 20036
Telephone: 202-331-9789
Toll Free: 800-982-2578
Fax: 202-872-0286
<http://www.wbonepub.com>

Request Reader Service #206

CP Graphics Library for Windows

Fast, dependable, royalty-free... and with the unique Routine Viewer your functions are just a few mouse clicks away.

- Bitmap and DIB functions for loading, saving, scaling, displaying, copying, capturing, printing...
- Multimedia and sound
- Transparent 'put' and sprite animation
- Clipboard routines
- User interface classes
- On-line help
- FREE fully documented source code \$250.00, delivery included.

ComputerPoint

71 Williamson Road, Para Hills 5096
Australia
Phone: (08)263 3623 Fax: (08)396 1477
email: cp@tne.net.au

Request Reader Service #172

Development CD-ROMs

Absolutely the World's largest, best indexed and most current collections of technical shareware CD-ROMs for professional programmers since 1984. Priced from \$25. Database (DBF) directories describe all products on CDs in detail, and include information on all commercial tools as well. Updated at least every three months.

We have CD-ROMs with 898 MS-Access products, 804 Assembler files, 2178 C/C++ files, 1084 C++ only, 1452 AutoCAD, 3314 Clipper, 566 Delphi, 2434 FoxPro, 2526 NetWare Utils, 592 Spreadsheet Utils, 1269 Paradox, 1322 Pascal, 618 Science, 328 VBDOS/QB, 1728 VBasic. Also Clarion and WinNT. All include 229,000 record database of PC products in DBF and Windows HLP formats.

EMS Professional Shareware
4505 Buckhurst Ct.; Olney, MD 20832
Voice:(301)924-3594 Fax:(301)963-2708
EMail:ems@wdn.com
<http://www.wdn.com/ems>

Request Reader Service #176

RINCON
RESEARCH
CORPORATION

Software Engineers
Interested in
Developing Software
Solutions for the Future

Rincon Research Corp. is seeking Software Engineers to create and design leading edge digital signal processing application frameworks for research and scientific prototyping in time-critical environments. We have immediate openings in the following areas:

- OOD/A methodologies with C++ on multi-threaded/tasking/processing platforms
- Graphical framework design (Motif, WIN 95) and design of user interface (2D/3D)
- Distributed processing design and analysis
- Hardware control and device drivers for UNIX, VMS and Windows NT

We offer competitive salaries and are committed to employee health and welfare by offering a comprehensive benefit package including relocation assistance.

Mail, Fax or E Mail your resume to Rincon Research Corp./Dept. HR, 101 N. Wilmet Rd., Ste. 310, Tucson, AZ 85711 / 520/745.1436 or resume@rincon.com. Principals Only. U. S. citizenship required

Request Reader Service #192

JOBS!

DICE is looking for Data Processing, Engineering and Technical Writing professionals to fill open positions for companies nationwide.

DICE is a FREE online job search service, providing detailed information about current contract and fulltime positions across the USA. Please contact by calling ANY of these access numbers, using your computer & 1200-9600 baud Modem, 8-N-1.

California	408-737-9339
Georgia	404-523-1341
Illinois	708-782-0960
Iowa	515-280-3423
Massachusetts	617-266-1080
New Jersey	201-242-4166
Texas	214-691-3420
Internet	telnet dice.com
Web	www.dice.com

DATA PROCESSING
INDEPENDENT
CONSULTANT'S
EXCHANGE



A Service of D&L Online, Inc. (515) 280-1144

Request Reader Service #173

COPY PROTECTION

INTERNET ACTIVATION

AZ-Tech Software, Inc.
Leaders in Software Security

EVERLOCK

SAFE-D

EVERKEY

Internet/CD-ROM

CPU Lock

HDD Lock

Serialization

Registration

Date & Execution

Limits

Win95 * 3.11 * DOS

Phone: (816) 776-2700

Fax: (816) 776-8398

E-Mail: sales@az-tech.com CC

I have been writing a database library for my own interest and exercise over the last several weeks. What interested me about this topic was the varied number and composition of fields that could be encountered at run time. After iterating through many data types I settled on an array of void pointers that I would cast with (CAST) or DYNAMIC_CAST<> at run time.

Here is a code sample of how I set it up:

```
void **fields;
for (i=0;i<NumberOfFields;i++)
    switch(FieldType[i]){
        case 'C':
            ((void) (field))[i] =
                new string;
        case 'N':
            ((void) (field))[i] =
                new int;
    ...
}
```

The compiler does not know how to increment when cast to void. (This is similar to Smalltalk's bag class). Since I know the order and the type of each field I could enjoy the convenience of accessing them as an array. The problem is it is an array composed of heterogeneous types.

I had also considered unions and a generic class. The union presented two problems: first, it wasted space; second, I wished to access the fields by the same name.

Another thought was a generic class that held all fields as characters. I felt this required too much conversion and I also would need a single function that returned different types depending on the type that the character sequence represented.

In short, I am looking for an array of pointers that allow me to access these elements, regardless of type, by the same name (ie. field[i].value).

I hope I have been clear here. Thank you very much for your time.

Sincerely,

Phill Moran

You're pretty close. Just define typenames for the various field types, then cast the void pointers appropriately, as in:

```
void **fields;
for (i=0;i<NumberOfFields;i++)
    switch(FieldType[i]){
        case 'C':
            *((Field_type_C *)field[i]) =
                new string;
        case 'N':
            *((Field_type_N *)field[i]) =
                new string;
    ...
}
```

— pjp



PROGRAMMERS... \$100k/yr at HOME!

My new book, "How To Make \$100,000 A Year And More Developing Low Budget Software Products", will reveal all the marketing tricks, strategies, and systems that have my business exploding with PROFITS!

Add your skills to my proven strategies and you have the PERFECT BUSINESS... Low overhead, part-time, home-based, huge margins, and UNLIMITED POTENTIAL.

CALL 800-364-4883

Call TODAY for a FREE special report!

ULTRA
Financial Systems
1633 Arrowhead Dr. Flower Mound, TX 75028

Fax: 214-724-0375

Compuserve: 71223,634

□ Request Reader Service #201 □

Now For Windows 95 Too!

Supports Visual C++, BC, 3.1, 95 & NT

NO ROYALTIES

FULL SOURCE

ALL for \$499.00

WAKE UP to the possibilities of advanced scientific graphics in your own software. Save years of programming time and effort.

Get full source code for advanced graphics without paying any royalties. Do not get into the situation of negotiating royalties after spending the time to learn and incorporate a graphics library into your program.

Get powerful and robust source code that is easy to understand and modify. Write to any windows device context, including windows, printers, icons, bitmaps and metafiles. Copy metafiles and bitmap graphics to the windows clipboard for export to other applications.

Get all the advanced graphics that you want: xy, error, Smith, bar and pie charts. Get solid or wireframe surfaces with color coded height and solid or labeled line contours, even contours on spheres.

Call (619) 632-9510 & Order It.

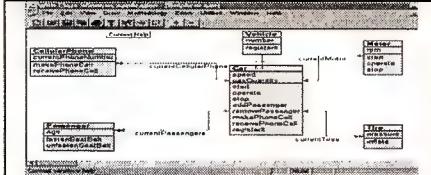
CHIRP Technical Services

P O Box 551, Del Mar, CA 92014

<http://www.chirp.com>

□ Request Reader Service #171 □

0-0 C++ Design Tool



- Customize Code Generation for State Diagrams in almost any Language - Ada, VHDL, C, Delphi, Pascal, Smalltalk, more
- Standard Scripts available for C, VHDL, C++, Pascal; Custom Reports
- Template driven code for State, Class and Object Interaction Diagrams
- Reverse Engineer C++, Delphi code

32, 16 bit **WITH CLASS** starts at \$295

Call for FREE Eval: **908.668.4779**

MicroGold Software, Inc.
76 Linda Lane Edison, NJ 08820
www.micropgold.com

□ Request Reader Service #185 □

Expert System for

C/C++ and Visual Basic

The Haley Enterprise, the leader in reasoning technology, offers **Eclipse**, **Rete++** and **VBXpert**, a family of embedable inference engines. Embed rule-based knowledge in your applications.

The first month is FREE!

The Haley Enterprise, Inc.
413 Orchard Street
Sewickley, PA 15143
TELE:(800)233-2622 FAX:(412)741-6457
E-Mail:Info@Haley.COM

(after which a permanent security code will be issued for \$199)

□ Request Reader Service #177 □

WROX PRESS

The Revolutionary Guide To MFC 4 Using Visual C++

Mike Blaszcak

ISBN 1-874416-82-3

\$49.95

800 Pages + CD



- Written by one of Microsoft's leading MFC developers
- Updated for Visual C++ 4.0 and MFC 4.0
- Explains the use of threads, exception handling and more
- Covers Windows NT 3.51 and Windows 95
- For more details visit Wrox Press Developer's Reference on the web at: <http://www.wrox.com>

How To Order

Phone 800-937-5557 Fax 800-PRI-ORDER

Visa MC Amex DS Quote code CL88 for free freight

For more information or a free color catalog of Wrox titles please call 800-USE-WROX or 312-465-3559.

□ Request Reader Service #207 □

C/C++ Users Journal — September 1996

C, C++ and Delphi Libraries for Windows and NT

G_FFT Class Library Extremely high speed Signal Processing routines, Real, Complex FFT, Correlation, Convolution, Amplitude, Power Spectrum, Filtering, Decimation, Data Smoothing, Windowing, sorting and more. **\$149**

G_FFT Professional: Includes G_FFT and supports Virtual Memory. Converts over 2 billion data points. **\$199**

OOPlot Class Library Plotting and charting, Linear, Log, Log-Log, Scatter, Bar, Pie, Origin setting, Scaling, Real Coordinates not integer, MDI support and much more. **\$149**

OOParser Class Library Expression and function evaluator with 1,2, & 3 variables and unlimited parameters. **\$79.**

Sigma Software, Inc.

15779 Columbia Pike, Suite 360
Burtonsville, MD 20866
Call or Fax your order: 301-549-3320
BBS 301-549-4161
Download Demo from our BBS.

□ Request Reader Service #198 □

Legacy C Complex C++

Easily Understand/Document Unfamiliar or Complex Code

- Wide range of views - control flow, data flow, class use, build dependencies, test coverage, performance...
- Knowledge-based exploration engine for data filtering
- Automatic design document generation, including html format
- 3D graphic visualization for rapid understanding
- Available on UNIX platforms
- Affordable price



Imagix

Free Trial Software
www.imagix.com
or call (503) 614-1100

□ Request Reader Service #178 □

Need to reach 45,000 C/C++ programmers with information about your products?

Advertise in:

C/C++ Users Journal™

Advanced Solutions for C/C++ Programmers

Call 913-841-1631 today
for information about
advertising opportunities in the
C/C++ Users Journal.

Advanced solutions for C/C++ developers.

breakout! marketing - Continental Europe +49 431-801740

Ed - East. Christine - Midwest. Julie - West.
913-838-7547 913-838-7546 913-838-7541

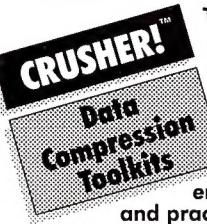
Looking for a ... Career Change ?

National Engineering Resources, Inc. can help you achieve your career objectives. We have many employment opportunities for computer professionals and C/C++ specialists.



Please send your resume to: NER, Inc.,
Attn. MS, 6200 Shingle Creek Parkway,
Suite 160, Brooklyn Center, MN 55430
Joblines: 800/665-7610 - Fax: 612/561-7675
nerinc@sprynet.com - <http://www.occ.com/ner/>

□ Request Reader Service #186 □



The best high performance, portable compression libraries for DOS, Windows, OS/2, Unix, Macintosh, embedded systems, and practically anything else, period.

FREE DEMO

DOS	\$249
Win16	\$299
Win32	\$299
OS/2	\$349
Unix	\$349
Macintosh	\$299

45-function API
Buffer compression
File compression
Disk spanning
Encryption
Self-extracting EXE's
V рХ/OCX controls
On-line help
Full source code

DC Micro Development

Tel 606-268-1559
Fax 606-266-0726

Call 1-800-775-1073

info@dcmicro.com

<http://www.dcmicro.com>

□ Request Reader Service #174 □

SPELL CHECKING & THESAURUS

■ FIRST CLASS ■ AFFORDABLE ■ ROYALTY-FREE
■ FAST - EASY - POWERFUL
■ U.S. ENGLISH THESAURUS TOOLKIT TOO!

DOWNLOAD OUR FREE DEMO !
<http://www.mv.com/ipusers/lexsaurus>

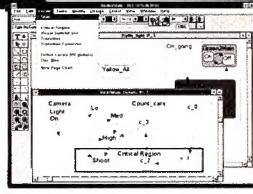
MULTI-LINGUAL	DOS
Danish	WIN 3.1
Dutch	WIN NT
French	WIN 95
Italian	OS/2
German	
Norwegian	
Spanish	
Swedish	
UK English	US English

LEXSAURUS SOFTWARE, INC.
427-3 Amherst St. Suite 303
Nashua, NH 03063
(603) 672-5941 • 1-800-570-1984 • Fax: (603) 672-5934
E-mail: info@lexsaurus.mv.com

□ Request Reader Service #181 □

Harel Statecharts

Java
C
C++
MFC



Delphi
Visual
Basic

- With **BetterState** you design & prototype visually using **Statecharts** (with Hierarchy, Concurrency, Critical Regions and more).
- User controlled Automatic code generation.
- Your code is always in sync with design specifications.
- Automatically generates multithreaded C++ and Java code.

URL: <http://www.r-active.com>

Phone: 408/252-2808
Fax: 408/777-8615
email: info@ractive.com

□ Request Reader Service #190 □

ProIndex™

ProIndex™

Full-Text Indexing and Retrieval
Development Toolkit!

- Perfect for CD-ROM applications and dynamic data management
- Unequaled indexing speed!
- Fast complex searches
- Handles wildcards, phrases, proximity and more!
- Use with C/C++, VB, Delphi, etc.
- LAN and multi-user support
- DOS, Windows, Windows NT, Windows 95, OS/2, NeXT, Unix, and Macintosh

Call today for a FREE demo and information!

InfoSphere
801-221-5902

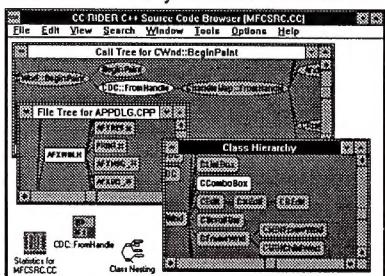
P.O. Box 225, Pleasant Grove, UT 84062

□ Request Reader Service #180 □

CC-RIDER

C and C++
BROWSING
FOR ANY EDITOR !

Source Code Analysis and Documentation



NEW v5.1 for Win 3.1, NT, DOS & OS/2

Western Wares

Free Demo (303) 327-4898

□ Request Reader Service #204 □

Attend the
WORLD RAIMA
USER CONFERENCE
Sept 23-24 in Strasbourg, France

"What makes Velocis the *fastest* database?"

You do! Velocis™ Database Server gives you speed-enhancing architectural choices and APIs, including SQL C-API, low-level C-API, C++ class libraries, and support for custom APIs. Typical relational client/server database products just can't keep up.

And there's more! Velocis supports both relational and pointer-based network model databases in any combination as well as processing on either side of the client/server equation.

The choices of multiple operating platforms, APIs, processing localities (client or server), and database models can be combined to satisfy the performance requirements of virtually any application, regardless of database size or complexity.

Call 1-800-ASK-4-RAIMA



For more information on what makes
Velocis the fastest Database, call for a copy of
the Bill of Materials Benchmark.

Velocis	Typical RDBMS	Features
✓	✓	ANSI SQL (DML, DDL, DCL)
✓	✓	Client/Server Architecture
✓	✓	RPC Mechanism
✓	✓	SQL C-API
✓		Administrative SQL C-API
✓		Low-Level C-API
✓		C++ Class Libraries
✓		Custom API
✓		Both SQL & C Stored Procedures
✓		C-based Triggers
✓		User Defined Functions
✓		Server Extensions
✓	✓	Relational Model
✓		Pointer-based Network Model
✓		Combined Database Model
✓	✓	Primary/Foreign Keys
✓	✓	Referential Integrity
✓	✓	Transaction Processing
✓	✓	Transaction Protection
✓	✓	Roll Forward Recovery
✓		Asynchronous Transactions

Test drive Velocis Database Server today at www.raima.com.

Raima Corporation, 1605 NW Sammamish Rd. #200, Issaquah, WA 98027 USA
Inside the U.S.: 1-800-275-4724, Outside the U.S.: 206-557-0200, Fax: 206-557-5200, Web: www.raima.com, Internet: sales@raima.com

Germany: 49 7022 9256 10 Singapore: 65 334 0061 Italy: 39 49 80 77 140 Argentina: 54 1 470 72 03 Spain: 34 3 225 39 95
Raima UK Ltd.: 44 1 273 819 292 Australia: 61 2 419 7177 Denmark: 45 44 88 99 00 Raima Benelux: 31 35 694 4738 Colombia: 57 1 62 10 070
France: 33 1 446 100 42 Sweden: 46 13 311 588 Baltic States: 372 631 4009 Finland: 358 0 8045 130 Japan: 81 3 5548 6926 Russia: 7 812 316 1965

Fast Database Engine

Now discover CodeBase 6.1 – The fastest and smallest database engine available with xBASE file compatibility!

Query a million records in just 0.97 of a second, or add ten thousand records in just 2.17 seconds.

Discover these valuable benefits of CodeBase:

- C, C++, Visual Basic and Delphi are all supported.
- FoxPro, Clipper and dBASE file compatible, and multi-user compatible
- Runs under Windows 95, 3.1, NT, DOS, and UNIX
- Tiny DLL loads quick and takes little memory
- Transaction processing and logging
- Client/Server configuration included
- Crash recovery is built-in
- Full featured report writer included
- Source code included
- Royalty Free—even with Client/Server

SQL Option Available

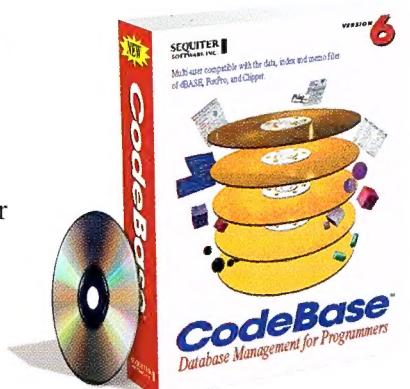
Now get the most database power at the lowest cost.



Award-Winning performance 5 years in a row!

"CodeBase gives ACT! the fast database access contact management users need"

- Michael Plasterer, Director of Development,
Symantec



The Database Engine for Programmers

FREE 30 Day Test Drive

Test drive the new CodeBase 6.1 for 30 days with your own code. No risk. No obligation. Order yours today.

Call: 403-437-2410

SEQUITER SOFTWARE INC. Fax: 403-436-2999
Email: info@sequiter.com
P.O. Box 575 Newmarket NH 03857-0575 Web site: www.sequiter.com

Request Reader Service #101



You can't find a better client SDK with these features on the market!

Over sixteen years of proven reliability and performance.

No one else supports over 30 platforms in this price range!

Q: What does it take to deploy a superior client/server application?

A: A SUPERIOR SERVER

START with the most advanced client-side SDK on the market: c-tree® Plus at \$895.

- Complete "C" Source code
- ROYALTY FREE (Client Side)
- Multiple supported protocols
- Fast, portable, reliable

- Powerful features like transaction processing
- Win95, NT, and Windows 3.1 ready

ADD a strong, multi-platform, industrial-strength Server that supports.

- File mirroring
- Heterogeneous networking
- Automatic disaster recovery

- Multi-threaded design
- Best price/performance available: from \$445- \$3745

RESULT? A solid, economical, easily deployable product that fits your needs.

- Portable
- Scalable
- Exceptional Performance

- Flexible
- Easy Server distribution
- Convenient OEM terms

c-tree Plus®

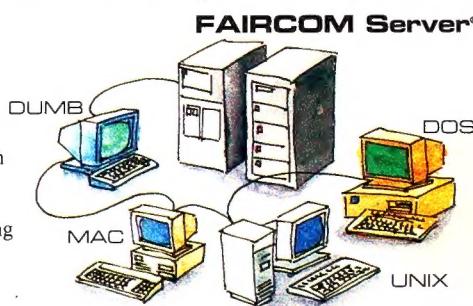
After 16 years, still the best in single-user, multi-user and client/server development. Now available for Windows 3.1, Win95, and Windows NT. c-tree Plus is still distributed in complete C source code and is known for its unparalleled flexibility, portability and royalty-free licensing policy. This licensing puts your development budget to work for you.

- Complete C Source
- Single/Multi User
- Client/Server (optional)
- Full ISAM functionality
- No Royalties
- Transaction Processing
- Fixed/Variable Length Records
- High Speed Data/Index Caching
- Batch Operations
- File Mirroring
- Multiple Contexts
- Unsurpassed Portability

FairCom Server®

Many major companies have turned to FairCom to implement their client/server systems. FairCom Servers adhere to the latest client/server design principles. FairCom Servers communicate using NETBIOS, SPX, TCP/IP, AppleTalk, shared memory, Message Queues or StreetTalk, depending upon the server platform. Most server platforms support more than one protocol, providing you the flexibility to choose how your client applications communicate to the server. For example, the FairCom OS/2 Server can communicate with shared memory, NETBIOS, SPX, and TCP/IP clients at the same time. No one else offers these features and portability for the price:

- Client/Server Model
- Transaction Processing
- Requires <2MB RAM
- Online Backup
- Disaster Recovery
- Rollback - Forward
- Anti-Deadlock Resolution
- Client-side "C" Source
- Multi-threading
- Heterogeneous networking
- File Mirroring
- OEM/Source Available



Heterogeneous TCP/IP Network

FOR YOUR NEXT PROJECT CALL FAIRCOM: YOU CAN'T FIND A BETTER HETEROGENEOUS CLIENT/SERVER SOLUTION!



FAIRCOM
CORPORATION

U.S.A. 4006 W. Broadway - Columbia, MO 65203-0100 • phone (573) 445-6833 fax (573) 445-9698

EUROPE Via Patrioti, 6-24021 Albino (BG) - ITALY • phone (035) 773-464 fax (035) 773-806

JAPAN IKEDA Bldg. #3,4f-112-5, Komei-chou - Tsu-city, MIE 514 Japan • phone (0592) 29-7504 fax (0592) 24-9723

Also inquire about these FairCom products:

d-tree™ r-tree® ODBC Driver

WWW Address: <http://www.faircom.com/>

(800)234-8180